REPORT DOCUMENTATION PAGE

AFRL-SR-BL-TR-98-
0414

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE	3. REPORT TYP
	1992	Final

4. TITLE AND SUBTITLE
A Psychophysical Investigation of Mechanisms Selective for the Spatial Frequency of Disparity Modulation in Random-Dot Stereograms

5. FUNDING NUMBERS

6. AUTHORS
Alan Blake Cobo-Lewis

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Wisconsin-Madison

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AFOSR/NI
110 Duncan Avenue, Room B-115
Bolling Air Force Base, DC 20332-8080

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION AVAILABILITY STATEMENT
Approved for Public Release

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*
See attached.

14. SUBJECT TERMS

15. NUMBER OF PAGES

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

Introduction

The visual system must perform two major functions. The first is resolving spatial layout, and the second is object recognition. Because we live in a three-dimensional (3D) world, these tasks must be performed in three dimensions. That is, the visual system must assess 3D spatial layout and must recognize 3D objects. There are two types of information sources that could be used to perform both functions. Cues of the first type arise directly from characteristics of in the environment. These include shape-from-shading, depth from texture gradients, and motion parallax due to movements of objects in the environment. Cues of the second type arise from characteristics of the observer. These include motion parallax due to movements of the observer through the environment, and binocular disparity due to the relative positions of the two eyes. This paper is concerned with how the visual system processes the last cue: binocular disparity. Specifically, it addresses the visual system's processing of binocular disparity across multiple spatial scales.

Spatial Layout and the Physiology of Stereopsis

Many neurophysiological data have been accumulated on primate stereopsis since Hubel and Wiesel (1970) first discovered neurons in macaque visual cortex selective for binocular disparity. Most of these data have served to identify basic types of disparity-selective cells (Poggio & Fischer, 1977; Poggio & Talbot, 1981; Poggio, 1984; Felleman & Van Essen, 1987; Poggio, Gonzalez, & Krause, 1988). These types include neurons narrowly tuned to zero disparity, neurons narrowly tuned to crossed or uncrossed disparities, and neurons broadly tuned to crossed or uncrossed disparities. In computer simulation, information pooling across these types of units has successfully described how humans' difference limen for disparity varies as a function of disparity (Lehky & Sejnowski, 1990). It is evident how these disparity-tuned neurons could underlie the

A PSYCHOPHYSICAL INVESTIGATION OF MECHANISMS SELECTIVE FOR THE SPATIAL FREQUENCY OF DISPARITY MODULATION IN RANDOM-DOT STEREOGRAMS

by

ALAN BLAKE COBO-LEWIS

A thesis submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

(Psychology)

at the

UNIVERSITY OF WISCONSIN-MADISON

1992

Table of Contents

Acknowledgment

This research was supported by a National Defense Science and Engineering Graduate Fellowship and by grant #135-3035 to Yei-Yu Yeh from the Wisconsin Alumni Research Foundation.

I thank my advisor, Yei-Yu Yeh, especially for her help in facilitating remote data collection. This project would not have been possible without this facilitation.

Special thanks are due my wife, Ceci, who made uncountably infinite sacrifices in order that this project be completed.

This work is dedicated to the memories of Rose Marie Farhy and Cecilia María Cobo.

Introduction

The visual system must perform two major functions. The first is resolving spatial layout, and the second is object recognition. Because we live in a three-dimensional (3D) world, these tasks must be performed in three dimensions. That is, the visual system must assess 3D spatial layout and must recognize 3D objects. There are two types of information sources that could be used to perform both functions. Cues of the first type arise directly from characteristics of in the environment. These include shape-from-shading, depth from texture gradients, and motion parallax due to movements of objects in the environment. Cues of the second type arise from characteristics of the observer. These include motion parallax due to movements of the observer through the environment, and binocular disparity due to the relative positions of the two eyes. This paper is concerned with how the visual system processes the last cue: binocular disparity. Specifically, it addresses the visual system's processing of binocular disparity across multiple spatial scales.

Spatial Layout and the Physiology of Stereopsis

Many neurophysiological data have been accumulated on primate stereopsis since Hubel and Wiesel (1970) first discovered neurons in macaque visual cortex selective for binocular disparity. Most of these data have served to identify basic types of disparity-selective cells (Poggio & Fischer, 1977; Poggio & Talbot, 1981; Poggio, 1984; Felleman & Van Essen, 1987; Poggio, Gonzalez, & Krause, 1988). These types include neurons narrowly tuned to zero disparity, neurons narrowly tuned to crossed or uncrossed disparities, and neurons broadly tuned to crossed or uncrossed disparities. In computer simulation, information pooling across these types of units has successfully described how humans' difference limen for disparity varies as a function of disparity (Lehky & Sejnowski, 1990). It is evident how these disparity-tuned neurons could underlie the

perception of spatial layout. Neurons selective for disparity would define a map of the depth at each location in the binocular image, and this depth map could directly serve the perception of spatial layout.

Shape Recognition and Stereopsis

If an approach to the perception of spatial layout is so straightforward, an approach to object recognition is more tortuous. An object cannot be recognized solely on the basis of its location in 3D space. What is important is its 3D *shape*. This is essentially a problem of pattern vision. What is required is a mechanism that extracts specific properties of the depth map for further analysis. Other properties of the depth map are irrelevant; for example, an object's 3D shape does not depend on its absolute depth. This extraction of form-in-depth, independent an object's specific depth, has been termed *hypercyclopean perception* (Tyler, 1975), after Julesz's (1960) term *cyclopean perception*, which refers to perception of image defined in binocular disparity, irrespective of the specific luminance information in each of the two monocular views. A cyclopean image is unitary, not dual, in spite of its arising from two separate monocular views. It is this unitary nature of the cyclopean image that gives it its name (cyclopean from the one-eyed cyclops). Just as cyclopean perception arises from abstracting patterns in the monocular views, hypercyclopean perception arises from abstracting patterns in the cyclopean view.

Pattern Vision and Channels

There have been myriad studies of the physiology of pattern vision, almost all of which have concentrated on patterns of luminance variation. An early extension of Hartline's (1938) concept of a receptive field (RF), which describes how a visual neuron is selective for the location in the visual field of a spot of light, was the consideration of excitatory and inhibitory areas within the RF. Thusly extended, an RF describes a neuron's

selectivity for certain distributions of light--in other words, for luminance patterns.

Neurons can be selective for more abstract patterns, as well. For example, neurons in cat visual cortex (Hubel & Wiesel, 1959, 1962) and monkey striate cortex (Schiller, Finlay, & Volman, 1976a; DeValois, Yung, & Hepler, 1982) are often selective for the orientation of a stimulus, independent of its location. It is of particular relevance to the present study that striate neurons are also often selective for the spatial frequency (SF) content of a luminance contrast pattern (DeValois, DeValois, Ready, and von Blanckensee, 1975; Schiller, Finlay, & Volkman, 1976b)

Different neurons can be selective for different subregions of the same dimension. For example, many monkey striate neurons are selective for SF, but the SF at which a neuron's sensitivity peaks varies over 4 to 5 octaves, even for neurons whose RFs lie at the same retinal eccentricity (DeValois, Albrecht, & Thorell, 1982). When different units are selective for different subranges of a dimension, it necessarily follows that these different subranges are processed at least partially separately. A mechanism, whatever be its structural implementation, that processes a specific sub-range of a dimension can be called a *channel.*

In the study of the perception and processing of luminance patterns, theories involving the coding of the SF of the luminance contrast distribution incident on the retinae by SF selective mechanisms have been applied to both neurophysiological and psychophysical results (for overviews, see Braddick, Campbell, & Atkinson, 1978; DeValois & DeValois, 1988). Analogous bandpass mechanisms, selective for the SF of binocular disparity modulation, have been postulated (Tyler, 1975, 1980, 1983; Tyler & Julesz, 1978; Schumer & Ganz, 1979), and it is these *disparity SF channels* that are the object of the present investigation.

A channel that is selective for SF can be viewed as a mechanism built around linear

filters. In the case of a disparity SF channel, the input to the channel is a scalar field representing disparity at each location in the image. The output from the channel is a function of the amount of energy the stimulus has in the band of SFs for which the channel is selective. If the channel is selective for SF only, and not phase, then the channel's SF filtering characteristics can be characterized solely by an amplitude spectrum in the SF domain. A linear-filter model of the channel would incorporate two filters with identical Fourier amplitude spectra, but which are in quadrature phase. The output of the channel would be a function of the sum of the energy passed by the two linear filters. Because of the natural applicability of the filter approach to SF channels, it is often convenient to speak of two stimuli processed by the same channel as being passed by the same filter.

Psychophysical Investigation of Disparity SF Channels

Though a single neuron can be viewed as a channel, it is also common to posit other (probably large, multi-neuron) mechanisms that act as channels. Though these channels have not been observed with currently available neurophysiological and anatomical techniques, their existence can be demonstrated via psychophysical techniques.

For disparity SF channels, a reasonable point of departure is to determine the selectivity of the visual system as a whole to disparity SF. A number of investigators have conducted this experiment (Tyler & Julesz, 1978; Schumer & Ganz, 1979; Rogers & Graham, 1982)--they traced out the function relating humans' sensitivity to disparity modulation in random-dot stereograms to the SF of disparity modulation. This function is analogous to the contrast sensitivity function (CSF) for luminance gratings. Tyler and Julesz reported peak sensitivity at about 0.3 cpd, and Rogers and Graham reported peak sensitivity for disparity modulation occurs at a spatial frequency of 0.2 to 0.4 cpd, while Schumer and Ganz reported peak sensitivity at 0.5 to 0.8 cpd. This is about an order of magnitude lower than the SF at which the luminance CSF peaks (Nachmias, 1967). In

Schumer and Ganz's experiment, sensitivity rolled off at about 5 to 6.5 dB/oct for higher SFs, up to 3 cpd (the highest frequency tested). Toward lower SFs, rolloff was more gradual.

Two qualitatively different types of model can explain a system's CSF: a single-channel model and a multi-channel model. In the single-channel model, the CSF is related to the tuning function of a single underlying filter. In the multi-channel model, there are several filters, each selective for a different subregion of the frequency axis; the CSF is related to the envelope of the sensitivities of these separate filters.

Psychophysical experiments can demonstrate the existence of multiple channels selective for different sub-ranges of a stimulus dimension. If stimulus A and stimulus B interact in their psychophysical effects, but stimulus A and stimulus C do not interact in the same way, then there must be some part of the visual system that processes both stimulus A and stimulus B, and which constitutes the site of interaction. Stimulus C, however, must not be processed at this site in the same way. In the language of channels, there exists a channel that processes stimulus A and stimulus B, but does not process stimulus C (Braddick et al., 1978). In the specific case of SF channels, their existence can be inferred from the selectivity of psychophysical effects to the spatial frequency properties of the image.

The effects of channels can be measured via several different psychophysical methods. These methods differ from one another in two fundamental ways (Braddick et al., 1978). First, different psychophysical methods can investigate different psychophysical interactions. An example of a specific interaction is that the processing of one stimulus may be impaired following prolonged viewing of another stimulus. A specific psychophysical method can determine the characteristics of a specific psychophysical interaction, but interactions investigated by different psychophysical

methods may show different characteristics. This would not invalidate any particular method, but would suggest that different sets of channels underlie the different interactions. However, if disparate methods yield similar ranges over which the interactions occur, it would permit the theoretical triumph of explaining a broad spectrum of results with a single underlying set of channels.

In addition to differing in the interaction under investigation, psychophysical methods for channel investigation can also differ in the psychophysical function under examination. Channels could operate only for specific psychophysical functions, but not for others. Or, different sets of channels could operate for different functions. For example, detection and discrimination represent different psychophysical functions that need not, *a priori*, be mediated by the same set of channels. However, as is the case for the role of channels in different interactions, when psychophysical methods investigating different psychophysical functions converge in yielding similar results, it adds theoretical power to the channel idea.

Sub-Threshold Summation. One common psychophysical method for demonstrating the existence of channels is *summation of sub-threshold stimuli*. In this method, the interaction is that if stimulus A and stimulus B, both of which are of sufficiently low amplitude that the visual system cannot readily detect them, are processed by the same channel, the output of which determines detectability, then when presented together, they might exceed the detection threshold. If, however, they are processed by independent channels, then the probability of detection would be the probability of detecting either stimulus independently of the other. Schumer and Ganz (1979) used this paradigm to find evidence for mechanisms selective for the spatial frequency of disparity modulation. They measured the threshold for detection of a sinusoidal disparity grating in the presence of a disparity grating of different SF whose amplitude was below detection

threshold. A single-channel model predicts that the presence of the sub-threshold grating would influence the threshold for detection of the test grating, and that the effect may depend on the relative phase of the two gratings, for harmonically related frequencies and gratings of finite extent (the latter includes all physically realizable stimuli). A multi-channel model predicts that the presence of the sub-threshold grating would only influence the threshold for detection of the test grating when they are closely spaced enough on the frequency axis to be passed by the same channel. The actual stimuli whose detectability Schumer and Ganz assessed were compound disparity gratings containing the first and third harmonics of a given fundamental spatial frequency. Whether the harmonics were in peaks-add or peaks-subtract phase, detectability was well predicted by a probability-summation model, which assumes that the compound is detected when either harmonic is detected independently. This indicates that the harmonics were passed by different channels. Because only first and third harmonics were used, the experiment cannot yield precise estimates of the SF selectivity of the channels.

The sub-threshold summation method is attractive for the investigation of SF channels. This attraction derives from the fact that any model of an SF channel must incorporate linear filters. Although a complete model of a SF channel likely needs to incorporate nonlinearities as well, it is the linear properties that are of primary interest, because SF filtering is a linear process. Thus, it is desirable to choose stimuli that the channel processes without substantial nonlinear distortion. The use of low amplitude stimuli serves to minimize the operation of commonly occurring physiological and psychophysical nonlinearities, such as compressive nonlinearities. However, using the low amplitudes that allow one to safely neglect these nonlinearities also means that stimuli are near or below absolute threshold, to which there are ecological objections. Namely, the visual system normally operates on stimuli well above threshold, and if near-threshold

stimuli are processed differently (e.g., if they escape nonlinear distortion), then experiments using these low amplitude stimuli could fail to yield insights into the processing of more common, and more important, stimuli. A related objection is that nonlinearities may represent a crucial aspect of visual processing. Carefully avoiding nonlinearities, or not at least searching to ascertain their existence or nonexistence, could therefore produce an incomplete picture of the visual system.

Selective Adaptation. Another common psychophysical method for demonstrating the existence of channels is *selective adaptation.* In this method, the interaction is that following prolonged exposure (adaptation) to stimulus A, sensitivity to stimulus B is reduced. If adaptation to stimulus A leaves sensitivity to stimulus C unaffected, then the adaptation is selective, and there must be a channel selective for stimuli A and B, but not C. The reduction in sensitivity following adaptation is generally understood as indicating neuronal fatigue. In addition to using sub-threshold summation, Schumer and Ganz (1979) also used selective adaptation to find evidence for mechanisms selective for the SF of disparity modulation. Their experiment was based on the threshold elevation technique used to investigate channels selective for luminance SF (Pantle & Sekuler, 1968; Blakemore & Campbell, 1969). Observers adapted to a suprathreshold sinusoidal disparity grating by slowly scanning their eyes across the peaks and valleys of the waveform, to eliminate local disparity aftereffects (Blakemore & Julesz, 1971). The adaptation gratings were presented at three times the threshold amplitude at that SF ($\approx$ 9.5 dB *re* threshold). The threshold disparity amplitude for detectability of subsequently presented disparity gratings was elevated by a factor up to about 1.7 ($\approx$ 4.8 dB). Threshold elevation was maximal for gratings presented at the same SF as the adapting grating, and threshold elevation fell with increasing distance on the frequency axis. At one octave above or below the adapting SF, elevation was about half maximum (6 dB down),

and at two octaves above or below the adapting SF, elevation was extinct.

Selective adaptation forfeits the advantage and escapes the disadvantage that sub-threshold summation realizes by using only low amplitude stimuli. However, the stimuli used by Schumer and Ganz (1979) were of relatively low amplitude; the stimuli thus might have escaped extensive nonlinear distortion by the visual system. One objection to using selective adaptation, though not to using the data generated, is pragmatic: because the method requires prolonged adaptation before detection trials begin, and frequent refreshment of adaptation while detection trials progress, the method is notoriously time consuming.

Selective adaptation is a tool with a specific purpose. Adaptation experiments ask what are the relatively long-term consequences of prolonged stimulation. This is a complementary question to that addressed by many other psychophysical methods: what are the immediate consequences of relatively delimited stimulation. Although these two questions may have similar answers, it is important to recognize that the questions do differ, and that they may have different answers.

Masking. A third common psychophysical method for investigating SF channels is *simultaneous masking*. In masking, as in selective adaptation, one stimulus (the masker) interacts with another (the signal) by reducing its detectability. In masking, however, rather than adapting to one stimulus and then viewing a subsequently presented briefer stimulus, the stimuli are instead of the same duration and are presented simultaneously. Tyler and Julesz (1978; Tyler, 1980, 1983) used a masking paradigm to investigate selectivity for disparity SF. They measured the detectability of sinusoidal disparity gratings in the presence of four-component pseudo-noises. The pseudo-noises were composed of four anharmonically spaced sinusoids spaced logarithmically across three quarters of an octave. They found extremely narrow tuning of about three quarters of an

octave.

Because rather intense stimuli are used as maskers, simultaneous masking, more than selective adaptation, runs the risk of nonlinear distortion. However, the extent to which such distortion actually occurs is an empirical question. Moreover, the nonlinearities, if they are well behaved and operate at appropriate sites in the visual system, may do no violence to the essential quality of SF selectivity. And, of course, because its stimuli are relatively intense, masking addresses the ecological objections to using stimuli near absolute threshold.

Adaptation in SF Identification. The three psychophysical methods detailed above investigate three different interactions--sub-threshold summation, selective adaptation, and simultaneous masking. However, all three methods investigate the same psychophysical function: detection of disparity gratings. All three methods thus demonstrate that there are mechanisms tuned to the SF of disparity modulation that function in grating detection, but the results do not necessitate that perception of SF is actually based on these mechanisms. Just as using medium- to high-intensity stimuli may allow masking to boast greater ecological validity than sub-threshold summation, using a psychophysical function more related to SF perception would allow a method to boast greater ecological validity than any method employing detection as its psychophysical function. Tyler (1975) conducted an adaptation experiment using just such a psychophysical function to address this issue. In his experiment, adaptation proceeded in a fashion similar to that of the paradigm of threshold elevation by selective adaptation, but the psychophysical function of interest was instead *SF identification*. His demonstration was based on Blakemore and Sutton's (1969) experiment with luminance modulation in which observers adapted to high contrast gratings and then reported the perceived SF of subsequently viewed gratings. The perceived SFs of gratings whose true SF was slightly higher than the adapting

frequency were elevated, and the perceived SFs of gratings whose true SF was slightly lower than the adapting frequency were lowered. The explanation is that perception of SF is based on the distribution of activity across SF channels. The perceived frequency is assumed to be the preferred SF of the most activated channel, or perhaps the SF corresponding to the centroid of the distribution of activity across channels. Via neural fatigue, the adaptation lowers the activity with which the adapted channel responds to subsequently presented gratings. This has the effect of shifting the centroid of activity away from the adapted channel whenever gratings of higher or lower SFs are presented. For example, if a grating is presented without adaptation, then it produces a distribution of activity across channels, and the centroid can be expected to occur near the grating's SF. However, if this same SF grating is presented following adaptation to a grating of higher SF, then activation of higher-SF channels will be depressed. Thus, there will be relatively more contribution to the activity distribution across SF channels from lower-SF channels than from higher-SF channels, which will shift the centroid toward lower SFs. Though he did not present quantitative results, Tyler (1975) demonstrated an analogous SF shift with square-wave disparity gratings. In addition, he presented data that demonstrated a shift in perceived orientation of disparity gratings following adaptation. The explanation for this phenomenon is that the perception of the orientation is based on mechanisms selective for orientation.

In addition to imbuing the idea of SF channels with ecological importance, the phenomenon of adaptation in SF identification also confers upon the idea greater theoretical power. Because SF identification represents a different psychophysical function than the previous three psychophysical methods investigated, adaptation in SF identification broadens the scope of the SF channel idea on a new dimension.

SF Discrimination. If a single set of channels underlies the phenomena revealed in

both detection experiments and identification experiments, then the results ought to be comparable. However, because the methods are based on different psychophysical functions, comparison of the results requires an elaboration of how the different psychophysical functions depend on the operation of the channels. This is a modeling enterprise, and comparison of the results thus tests the specific model elaborated.

Once a model is developed with the results of one of the above psychophysical methods, additional predictions from the model can be generated and tested. If the model is generally applicable, then it can even make predictions for the results of experiments not specifically designed to reveal the properties of channels. If such a model successfully predicts the results of such an experiment, then the SF channel idea demonstrates great theoretical power, for the simple idea is shown to account for results in truly disparate realms.

A pertinent example is provided by the experiments of Wilson and colleagues. Having used a masking paradigm to develop an elaborate model of luminance SF channels, including the shape of their filters' amplitude spectra (Wilson, McFarlane, & Phillips, 1983), they expanded the model and tested whether it could also account for the results of an experiment measuring the difference limen for luminance SF (Wilson & Gelb, 1984). The expansion was based on the insight that for a system equipped with SF channels, the task of SF discrimination could be recast into a task of intensity discrimination within a channel. Specifically, when presented with a sinusoidal luminance grating, an individual channel is stimulated to the extent that its filter passes energy at the grating's SF. When presented with another grating of similar, but not identical, SF the same channel is differentially stimulated to the extent that its filter passes a different amount of energy at the second grating's SF. If SF discrimination is based on the output of the SF channels, then gratings of SFs sufficiently different that they yield discriminable outputs of an SF

channel ought to be highly discriminable from one another, whereas gratings of SFs so similar that they do not yield discriminable outputs for any SF channel ought to be indiscriminable from one another.

Bandwidth of Disparity SF Channels

A fundamental question about disparity SF channels is: How sharply are they tuned for SF? In other words, what are the channels' bandwidths?

Bandwidth Estimates in the Literature. Because it used only first and third harmonics, Schumer and Ganz's (1979) sub-threshold summation experiment cannot yield a precise estimate of the bandwidths of the SF channels. However, it does verify that there exists a channel with a bandwidth less than the bandwidth of the stereovisual system as a whole, and it upper bounds the channel's half-bandwidth at three octaves. Tyler's (1975) SF identification demonstration, because it was not an experiment, also cannot provide an estimate of the bandwidths of the SF channels.

More estimates of the bandwidths of mechanisms selective for the SF of disparity modulation are available. They come from two sources, and the estimates conflict. Results obtained by measuring threshold elevation following adaptation indicate bandwidths of two to three octaves (Schumer & Ganz, 1979). On the other hand, results obtained by measuring threshold elevation produced by simultaneous masking indicate bandwidths of less than one octave (Tyler & Julesz, 1978; Tyler, 1983).

Explanations for the Conflict in Bandwidth Estimates. There are several reasons why these two methods could yield different estimates. First, two different mechanisms could underlie the phenomena investigated by the two different methods. Along these lines of reasoning, Tyler (1983) suggested that the narrowly tuned mechanisms that masking reveals might not be adaptable.

Alternatively, it is possible that a solitary mechanism does indeed underlie both

phenomena, but that the bandwidths of the two phenomena differ because of the intervention of distinct processes. It may be that the same channels selective for the SF of disparity modulation are affected both by adaptation and by masking, which would mean that a solitary mechanism underlies both phenomena; however, adaptation and simultaneous masking could have different effects on those channels (Graham, 1989). It is a common assumption that masking elevates threshold by adding noise to a channel's input, without changing the channel's fundamental properties. That is likely not the mechanism by which adaptation elevates threshold. The effects in adaptation experiments may involve relatively long-term changes in a channel's actual properties. For example, bleaching of pigments in retinal receptors is a long-term change in properties that is often invoked to explain adaptation effects in color experiments. Furthermore, even though there is evidence that masking is based on adding noise to a linear channel (see Results from Experiment 2, below), there is no guarantee that adaptation has a linear characteristic. This would distort the apparent shape of the channel, and hence its measured bandwidth (Tyler, 1983).

There is another possible explanation for why selective adaptation yielded such wide bandwidth estimates. It relies on Dealy and Tolhurst's (1974) conclusion in their study of channels selective for the SF of luminance modulation, that adaptation may be an aftereffect not of a channel's excitation, but rather of its inhibition. They based this conclusion on the results of an experiment in which they presented adaptation gratings below the threshold of a target channel. Presumably, because they were sub-threshold, these gratings did not excite the excitatory region of the channel's receptive field. Nevertheless, adaptation to these gratings elevated the channel's threshold. They concluded that the bandwidth of a channel's inhibition may be wider than its bandwidth for excitation. This conclusion implies that adaptation is not a processes of channel fatigue,

but rather of mechanism activation. By activating an inhibitory region, adaptation elevates thresholds. However, this explanation relies on the assumption that the inhibitory region is more activated by the adaptation than the excitatory region is. Otherwise, it would predict that threshold at the adapting frequency would be lowered rather than raised maximally. This model could explain why Schumer and Ganz obtained such wide bandwidths. However, it has been argued that the wide bandwidth of inhibition should also apply to simultaneous masking (Braddick, Campbell, & Atkinson, 1978). This explanation therefore predicts agreement, rather than disagreement, between the bandwidths obtained with adaptation and masking (Braddick, Campbell, & Atkinson, 1978).

Off-frequency Viewing. There is another possible explanation for why the masking curves are so narrow. Tyler and Julesz (1978; Tyler, 1983) used pseudo-noise maskers analogous to bandpass noises. When bandpass maskers are used, the width of a masking curve can be narrower than the tuning of the filters underlying the frequency selectivity. A system so equipped with multiple filters, and that implements an ideal detection strategy, bases the detection decision on the activation of the filter with the highest signal-to-noise ratio (SNR). When the masker is centered on the frequency axis about the signal, the filter with the highest SNR is the one centered about the signal (Figure 1, upper left). When the masker is moved slightly away in frequency from the signal, the SNR in this filter increases (Figure 1, lower left). However, if the filters are relatively flat near their peaks relative to their skirts, an additional gain in detectability is achieved if detection is based on the output of another filter centered about a frequency on the opposite side of the signal from the masker (Figure 1, right). This gain is achieved because, although this second filter attenuates the signal somewhat more than the filter centered about the signal, it attenuates the noise more, and it is the signal-to-noise *ratio*

that determines detectability. This strategy of basing detection on the output of a filter centered away from the frequency of the signal can be termed *off-frequency viewing*, and its consequence is that the masking curve is narrower than the tuning curve of the filter centered about the signal. To eliminate the same phenomenon in the measurement of the tuning of the auditory filter, a notched noise masker was introduced (Patterson & Nimmo-Smith, 1980; O'Loughlin & Moore, 1981). Such a masker is the sum of two noises with non-overlapping passbands, so that the noise's spectrum contains a notch. When a notched noise is constructed as the sum of a high-pass and a low-pass noise, it is equivalent to a bandstop noise. In a notched noise masking experiment, the signal is placed in the center of the notch. The tuning is measured by varying the width of the notch. By keeping the signal centered in the notch, one prevents off-frequency viewing, because the filter with the highest signal-to-noise ratio remains the one centered about the frequency of the signal.

Blake and Holopigian (1985), used a similar procedure to eliminate off-channel viewing in visual psychophysics. They investigated the detectability of luminance gratings masked by other gratings of the same SF but different orientations. Their analog for a bandpass masker was a single grating with an orientation slightly clockwise, say, from the signal's orientation. If the observer's response were based on the activity in channels tuned counterclockwise from the signal's orientation, then it would constitute off-channel viewing. Their analog for a notched masker was two gratings whose respective orientations were clockwise and counterclockwise from the signal's orientation. Because channels tuned both clockwise and counterclockwise from the signal orientation were stimulated by maskers, off-channel viewing would not be expected. The results were that the falloff of masking was indeed more gradual with the notched-noise analog than with the bandpass analog, which supports the existence of off-channel viewing.

In an unpublished study of the masking of disparity gratings by pseudo-noises, Tyler used two-band pseudo-noises analogous to notched noises. He compared the masking curves yielded by pseudo-noises composed of four anharmonic sinusoids logarithmically spaced over three quarters of an octave to the masking curves yielded by a notched noise analog, wherein two pseudo-noise components were placed at frequencies higher than the grating to be detected and two pseudo-noise components were placed at frequencies lower than the grating to be detected. Surprisingly, the masking curve yielded by "notched" pseudo-noise was actually somewhat narrower than the masking curve yielded by the "bandpass" pseudo-noise (Tyler, personal communication).

Tyler (1983) measured the amount of masking for several placements of the sinusoid in the "passband" of the pseudo-noise. He found that masking was constant as long as the signal was in the "passband". This is an odd finding, because when the signal is at the edge of the noise, a filter centered about the signal would reject most of the energy in the pseudo-noise, whereas when the signal is in the middle of the noise, a filter centered about the signal would pass most of the energy in the pseudo-noise. The presence of off-frequency viewing would only accentuate this effect. Thus, masking ought to be greater when the signal is in the middle of the pseudo-noise than it is when the signal is at the edge of the pseudo-noise.

Continuum Versus Small Number of Channels. One explanation for why masking is constant as long as the signal is within the "passband" of the pseudo-noise is that rather than having a continuum of filters, as occurs in the auditory domain with the critical band, there is instead of small number of channels tuned to the SF of disparity modulation. This is an issue common in the investigation of channels selective for the SF of luminance modulation (e.g., Wilson & Bergen, 1979).

There would be clear evidence for the existence of a small number of channels if

the masking curves for signals at a variety of SFs all peaked at just a few locations on the SF axis. Those locations would be the frequencies where the channels show peak sensitivity. This is because a masker at a channel's peak frequency would be the most effective masker for any signal to which that channel is the most sensitive, regardless of the signal's SF. What, then, if two channels were equally sensitive to a signal, such as might occur when a signal is placed midway between two filters of a bank discretely spaced filters? Detection of the signal would be mediated by probability summation across both filters. It is conceivable that masking would be maximal when either one of the two filters is effectively flooded with noise. This would produce a bimodal masking curve, with peaks at the center frequencies of the filters and a dip at the signal frequency. However, the bimodal masking curve need not obtain. If the shape and overlap of the filters are "chosen" carefully, then the most effective masker might be one that makes the "best" compromise between flooding both channels with noise. That noise could well be centered at the signal's frequency, for that frequency represents an obvious compromise between the two channels' peak frequencies. The possibilities for such a situation are multiplied if one considers that strict probability summation is but one of an infinity of possible rules for combining information across channels. Thus, although there would be clear evidence that there were only a few channels if the masking curves for many different signals peaked at just a few SFs, the lack of such a small number of peaks would not resolve the issue of whether there were only a few channels or whether there were a continuum of channels.

However, there is another way in which to obtain evidence for a discrete number of channels. Off-frequency viewing relies on the observer's detection process's having access to a tunable filtering process, so that the detection process can monitor SFs slightly different from the signal SF. If there is a small number of channels, then tuning is severely

limited. Specifically, if the signal were placed at a frequency midway between two channels, then opportunity for off-frequency viewing would be ample, because the observer's effective filter could be any one of the uncountably infinite linear combinations of the two nearest channels--the filtering process could be effectively be detuned away from the SF of the masker. However, if the signal were placed at a frequency where a channel is at peak sensitivity, then off-frequency viewing would fail to lower threshold, because the nearest available channel would be tuned to an SF distant from that of the signal (if the nearest channel were very close, then the system would be indistinguishable from one employing a finely tunable filter or a continuum of filters). Such a filter would be of little use, because it would attenuate the signal drastically, whereas the utility of off-frequency viewing relies on the signal's being only slightly attenuated. Thus, if there is a small number of channels, then as signal SF is swept across a bank of a small number of filters, the difference between notched masking curves and bandpass masking curves would successively disappear, reappear, and disappear again, as the signal SF fell at a filter peak, between filter peaks, and at the next filter peak. On the other hand, if there is a continuum of filters, then the difference between notched masking curves and bandpass masking curves would persist at all SFs.

Summary of Issues

A set of related issues has presented itself in the above discussion of disparity SF channels. A fundamental question is: What are the bandwidths of the channels? Addressing this question entails addressing the ancillary question: What is the extent of off-frequency viewing? A more specific question, Does off-frequency viewing disappear at some SFs?, is related to the fundamental question: Are there a small number of channels, or is there a continuum? A final question is: Is the idea of SF channels powerful enough to account for data in disparate realms?

Approach of Current Study

The goal of the current research is to deepen understanding of the visual system's disparity SF selectivity, especially by addressing the questions listed above. The approach adopted is to develop a rigorous quantitative model of the processing of a stimulus according to its SF characteristics. This approach has numerous qualities to recommend it.

For one, a rigorous model yields detailed predictions to confront data. By quantifying notions such as SF channels, a rigorous model can highlight specific areas where theory succeeds and other areas where theory fails.

To date, all quantitative accounts of the properties of SF channels have been largely data-centered. For example, instead of reporting bandwidths of the underlying mechanisms, investigators have instead reported bandwidths of the curves fitting the data. Although the data themselves must never be ignored, research must move beyond data description to data explanation. The bandwidths of SF channels is at least as important a measurement as the width of masking curves.

A related benefit of a modeling approach is that it forces one to rigorously derive the consequences of even a simple model. In the absence of this analysis, it is unknown to what extent a simple model can account for complex phenomena. Naive interpretations of the data might be proven erroneous if the consequences of the implicit model were derived. For example, if the visual system operates on the output of channels as an ideal observer, then a naive interpretation of bandpass noise masking curves would yield a spuriously low estimate of channel bandwidth, because it would ignore how off-frequency viewing leads to masking curves narrower than the tuning of the channels.

Another benefit of quantitative modeling is that many questions about SF channels can be asked instead of the model. Among the questions that such a model can answer is

whether there are a discrete number of SF channels or whether there is instead of continuum of filters (cf., the critical band in psychoacoustics). As mentioned above, such a model can also answer the question of what are the bandwidths of the SF-selective mechanisms, and it can address the question of to what extent are the different bandwidth estimates derived from different psychophysical paradigms due to coarse methods of analysis. However, the applicability of such a model is potentially broader than just answering these questions. Armed with a quantitative model of the selectivity for the SF of disparity modulation, one can predict manifold data, such as discriminability of the SF of disparity modulation for various SFs (shape discrimination), detectability of complex 3D surfaces (shape detection), etc.

The quantitative model developed herein is based on the results of a battery of experiments, described below, that measured detectability of unmasked and simultaneously masked sinusoidal disparity gratings. The immediate goals of the first five experiments are to use the data to develop a quantitative model to answer questions about the SF selectivity of the visual system and the extent of off-frequency viewing. A subsequent experiment that measured SF discriminability as a function of SF was used to test the model in the domain of a different psychometric function.

General Method

Subjects

Three observers served as subjects. Subject AC, a 25-year-old male, is the author. The other two subjects, AG, a 21-year-old female, and DR, a 25-year-old male, were naive to the hypotheses under test. All three have normal or corrected vision and good stereoacuity.

Apparatus

Stimulus presentation and data collection were performed by a 33 MHz 80386/387

PC with stereogoggle and TMS 34020 graphics boards supplied by Vision Research Graphics, Inc., Durham, NH. The stimuli were displayed on a Nanao 9080i 16-inch color monitor, operating in monochrome mode, with 1024 x 512 pixels (horizontal x vertical), refreshed at a rate of 119 Hz (noninterlaced). At the viewing distance of 75 cm, the 28 x 20 cm screen subtended 21.1 x 15.2 deg of visual angle. However, to avoid crosstalk between the two eyes' images, only the top 450 pixel rows were used. Thus, the stimulus images subtended 21.1 x 13.4 deg of visual angle. The monitor was operated in black-and-white mode, and luminance resolution of each pixel was approximately 15 bits. Stereo presentation was achieved by time-multiplexing images intended for the left and right eyes. The subject wore goggles equipped with liquid crystal shutters whose operation was synchronized with the display operation, so that when the monitor was displaying one eye's image, the goggles were transparent for the corresponding eye and opaque for the noncorresponding eye.

Stimuli

All stimuli were static random-dot stereograms (RDSs) displayed for 100 ms. If each element of the RDS were a single pixel, then disparity resolution would be limited by the width of a pixel. However, each element of the RDS was instead a horizontal *pair* of pixels. This allowed much finer control of the perceived horizontal position of an RDS element, because the luminance of each pixel in a pair could be controlled independently, and thus the luminance centroid could be shifted with a resolution well below the pixel width. Such sub-pixel resolution was limited only by the number of entries (256) in the gamma-corrected pseudo-color lookup table. At the 75 cm viewing distance, horizontal resolution of the luminance centroid in a pixel pair was nominally 0.2 sec.

This method of RDS generation differs from that introduced by Julesz (1960). In his method, all disparities must be an integral number of pixels, which precludes sub-pixel

resolution. Instead, in the present method, disparity is first defined as a function of vertical position, so that disparity is constant across horizontal position. This defines a horizontal disparity grating (one whose periodic energy is concentrated on the vertical SF axis). Five thousand x, y positions are selected randomly and independently for RDS element placement, and for each position, a two-pixel RDS element is plotted in the left eye's display and another in the right eye's display, at the appropriate disparity. Because disparity is not a function of x position, no monocular dot-density cues are introduced.

Sixty-four samples of each noise profile were created offline by C programs that generated gaussian noise and then filtered it appropriately via digital Fourier transforms. Sinusoidal disparity profiles were generated online, and the signal and noise profiles were added online as well. The phase of each sinusoid was determined randomly for each interval in each trial independently, as was the selection of the RDS element placements.

Psychophysical Procedure

Data collection proceeded in sessions of approximately ninety minutes each. Several experimental runs were conducted in each session. All data were collected by two-interval forced choice adaptive staircases. Conditions were blocked, so that during an experimental run, the SF of the signal and the long-term spectrum of any masker present was constant. This was to avoid uncertainty effects, which would otherwise tend to discourage ideal observing strategies. This is important, because off-frequency viewing constitutes one such strategy sensitive to uncertainty. Moreover, in order to maximize peformance, and consequently to maximize the chance of any off-frequency viewing to be detectable, auditory feedback was given following all trials on which the observer responded incorrectly. During each run, two independent staircases were randomly interleaved. Each threshold measurement derived from fitting psychometric functions to the data generated from at least two such runs.

Detection Experiments

In this first set of experiments, thresholds were measured for the detection of sinusoidal disparity gratings in the presence of notched-noise maskers, in the presence of bandpass-noise maskers, and in the absence of maskers.

Two differences between the maskers in this study and Tyler's are particularly noteworthy. First of all, in his studies, all maskers were four-component pseudo-noises. In the present study, all maskers were true noises, produced by digitally filtering large vectors of computer-generated gaussian noise. The second important modification is that the noises used herein were narrowband. The motivation for using such stimuli is that to probe a channel's structure, it is desirable to use tools that are finer than the channel under study. Because the channels may be narrowly tuned, narrowband noises are befitting.

For all measurements, signals ranged from 0.2 cpd to 1.6 cpd, in half-octave steps. The half-octave spacing was chosen so that if there is a small number of channels, then some signal placements would occur near a channel's peak SF and some signal placements would occur far away from any channel's peak SF.

The initial signal intensity for each staircase was about 29 dB *re* 0.1 arcmin RMS (amplitude of 4 arcmin). During a staircase's practice trials, which continued until six turnarounds in that staircase were completed, signal intensity was decreased by 3 dB following correct responses and increased by 12 dB following incorrect responses. Following the practice trials, signal intensity was decreased by 1 dB following correct responses and increased by 4 dB following incorrect responses. Each staircase continued after practice until ten additional turnarounds in that staircase were completed.

Experiment 1. Threshold as a Function of Masker Level

Subjects. AC and AG served as subjects in all this experiment's conditions. Additional data in two conditions were obtained with DR as a subject.

Design. In masking studies, it cannot be taken for granted that the mechanisms behave in a linear fashion. For that reason, Experiment 1 measured the amount of masking of a sinusoidal signal by a bandpass noise centered about that signal, for a wide range of noise intensities. For noises intense enough so that the noise in a channel is dominated by the masker noise rather than by the channel's intrinsic noise, a linear channel ought to yield a threshold proportional to the intensity of the masker. If such proportionality is observed, then the falloff in masking in the notched-noise and bandpass noise experiments below (Experiments 1-2) can be taken as directly indicating the attenuation due to SF filtering by the visual system.

Maskers in this experiment were 1/3-octave band noises centered about the signal frequency. Intensities ranged from about 19 to about 37 dB *re* 0.1 min (0.92 to 7.35 min RMS, respectively). For each subject, data-dense (i.e., measurement about every 3 dB) plots of masking as a function of signal level were be obtained for two SFs, and sparser (i.e., measurement about every 6 dB) plots were obtained for the other five SFs.

Results and Discussion. Initial threshold estimates were taken as the mean of the last 8 or 10 turnarounds in each staircase. These thresholds are displayed in Figures 2-4.

If channels behave linearly, then signal intensity at threshold ought to be proportional to the total intensity of noise in the channel. Furthermore, if the maskers are intense enough to dominate the noise in the channels, then signal intensity at threshold ought to be proportional to masker intensity. In a dB-dB threshold-versus-intensity plot, this proportionality exhibits itself as threshold rising with masker intensity with a slope of 1. The broken lines in Figures 2-4 display the best fitting lines with a slope of 1. It can be seen that these unity-slope lines fit the data well. Table 1 presents the proportions of variance accounted for by the unity-slope lines. These proportions are uniformly high.

Because a finite number of noise samples was used in each condition, the average

noise intensity the subjects saw could deviate from the expected noise intensity (Figures 2-4 are based on the average noise intensities the subjects saw). Although the mean absolute deviation across all maskers used in Experiments 2-4 was only 0.33 dB, it was in one case as large as 2.34 dB. However, because the data in Figures 2-4 demonstrate so clearly that threshold is proportional to masker intensity, correcting for a deviation is simply a matter of adding a dB offset to a threshold that is equal in magnitude and opposite in sign to the deviation of noise intensity in that condition. It is important to apply this correction to ensure that thresholds obtained in conditions where the maskers were of equal expected intensity but had different spectra be directly comparable. Henceforth, all data are thusly corrected.

Another feature of the data is particularly notable--masked threshold does not depend on the SF at which threshold was measured. To demonstrate this, Figure 5 plots the *y*-intercept of the unity-slope lines from Figures 2-4. This independence is important, because for stimuli of constant spatial extent, an ideal observer's threshold would be lower for the high-SF stimuli than for the low-SF stimuli. This is because noise is stochastic. Any given stimulus contains just a sample of the noise. If performance is based on the intensity of the signal-plus-noise stimulus relative to the noise-alone stimulus, then performance will improve as the observer has available a better estimate of the true long-term intensity in each interval. As the spatial extent of a stimulus increases, the variability in the sampling distribution of intensity decreases, leading to better estimates of intensity, and lower thresholds. For noises of constant octave bandwidth, the variability in the sampling distribution is held constant only if stimulus extent is scaled with SF. However, in the present experiment, all stimuli were of constant spatial extent. Thus, the variability in the sampling distribution was greater the lower the SF. More quantitatively, we can define the *coherence length,* the distance beyond which extended samples of the noise are

approximately independent, as

$$y_{\mathrm{coh}} = \frac{1}{f_g(c - 1/c)}, \tag{1}$$

where f_g is the signal frequency and center frequency of the noise, and, for 1/3-octave noise, $c - 1/c \approx 0.23$. Eq. 1 is derived in Appendix A. In the 0.2 cpd condition, the stimulus extends only over about 0.62 of a coherence length, or over about two-and-a-half cycles of the signal. The empirical result that performance is the same in the 0.2 cpd condition as it is in all the other conditions, wherein the stimulus extends over several coherence lengths, means that the human observers are never using more than 0.62 coherence lengths' worth of information. In effect, they summed information over only a limited spatial extent.

Given that human observers do not sum information over more than 0.62 of a 1/3-octave noise's coherence length, one can model the detection process and thereby model the entire psychometric function. The advantage of having such a model in hand is that, instead of basing threshold estimates on just the mean of the turnarounds in the staircases, one can fit all the raw data (proportion correct as a function of stimulus intensity) with a model of the psychometric function. The advantage of taking such an approach is that it extracts more information from the data than would be extracted by just using the mean of the turnarounds.

The model, fully derived in Appendix B, is a near-ideal observer of stimulus energy. It chooses as its guess for the signal-plus-noise interval that interval with a higher energy content. Because the noises in the present experiment are narrowband, it is assumed that any SF filtering conducted by human observers leaves them largely unattenuated. For this reason, the near-ideal observer does not filter these stimuli.

For each subject in each condition, raw data were pooled from all runs to yield

plots of frequency-of-correct-response versus signal intensity. This psychometric function was fitted according to the maximum-likelihood criterion with the near-ideal observer's psychometric function, which is identical to a Quick (1974) psychometric function,

$$\begin{aligned} \text{P(correct response)} &= 1-\frac{1}{2}2^{-S/(2N\ln 2)} \\ &\approx 1-\frac{1}{2}2^{-0.72S/N} \end{aligned} \quad , \qquad (2)$$

where S is the signal intensity, and N is a free parameter representing the total intensity of noise, both internal and external to the observer, that limits detectability (for absolute detection, there is only internal noise). Threshold was taken as the signal intensity that yields 75% correct. Figures 6-8 display the 75% thresholds from the raw data thus fit, along with the predictions of the near-ideal observer model. There is no plot for AG's 0.2828 cpd signal condition, because it was decided to retain data on individual trials after AG had already completed most of this condition. Averaged across all data points, AC's performance exceeds near-ideal performance by about 1.4 dB, AG's performance falls short of near-ideal performance by about 1.8 dB, and DR's performance exceeds near-ideal performance by about 3.2 dB. There are at least four possible explanations for how two of three subjects could perform so extraordinarily well as to apparently exceed near-ideal performance. The first explanation is that the human stereovisual system may filter the narrowband maskers so that their effective intensity is attenuated slightly. Because the near-ideal observer performs no filtering of the stimuli in this experiment, threshold estimates could be lower than near-ideal threshold. The second explanation is that the maximum-likelihood estimator of N in Eq. 2 may be biased downward, as maximum-likelihood estimators can be. For example, it is well known that the maximum-likelihood estimator of population variance is biased downward by the factor $(n-1)/n$, where n is

the sample size. This would lead to threshold estimates that are also biased downward. The third explanation is that human performance does not actually exceed near-ideal performance, but the results suggest otherwise because of sampling error. The last explanation is that human performance actually is extremely close to ideal, and thus truly exceeds near-ideal. Whatever the explanation, the discrepancy is small. The near-ideal model provides a good fit to the human data.

Experiment 2. Notched Noise

Subjects. AC and AG served as subjects in all this experiment's conditions. Additional data in two conditions were obtained with DR as a subject.

Design. In Experiment 1, thresholds were measured for the detection of sinusoidal disparity modulations in the presence of bimodal maskers of constant intensity. The maskers each had a 1/3-octave passband at an SF below that of the signal and another 1/3-octave passband at an SF above that of the signal. Signals were centered logarithmically in the notch between the passbands. By varying the width of the notch, one can measure the SF selectivity of the underlying channels in the absence of off-frequency viewing. For each of the seven SFs of the signal, notches of four widths were be used--passbands were centered at $\pm 1/2$, ± 1, $\pm 3/2$, and ± 2 octaves relative to the signal. For 1/3 octave passbands, this corresponds to notch widths of 1/6, 2/3, 7/6, and 5/6 octaves, respectively. The noise intensity in the two passbands were equal. The total noise intensity, 37.3 dB *re* 0.1 min arc, was the same as one level used in the measurement of threshold as a function of masker level (Experiment 2), so that by combining the present data with data from Experiment 1, one has data for five notch widths, the fifth being a degenerate "zero notch width" from Experiment 1.

Results and Discussion. Subject AG produced unstable data for the 0.8 cpd signal in notched noise with passbands centered $\pm 1/2$ oct *re* signal SF. That condition was re-

run with the masker at half-amplitude (6 dB down). The threshold reported below for that condition represents the half-amplitude result, plus 6 dB to make it directly comparable to the other thresholds. This procedure is justified by the unity-slope threshold-versus-intensity plots described in Experiment 1. The problem of unstable data is considered in more detail below, in Experiment 3, Results and Discussion.

The raw data were fit with the psychometric function displayed in Eq. 2. The rationale is that the noise parameter, *N*, which quantifies the effective intensity of the noise, reflects to what extent the human observer's SF filtering attenuated the noise. In conditions in which the noise was greatly attenuated, *N* will take on a low value, and in conditions in which the noise was not greatly attenuated, *N* will take on a high value, near to the value of the masker's full intensity. The squares in Figures 9-11 display the 75% thresholds thus obtained. The prominent feature of the data is that as notch width increases, threshold generally decreases. This is compatible with there being a continuum (or large number) of SF channels.

Experiment 3. Bandpass Noise

Subjects. AC and AG served as subjects in all this experiment's conditions. Additional data in some conditions were obtained with DR as a subject.

Design. In Experiment 3, thresholds were measured for the detection of sinusoidal disparity modulations in the presence of bandpass maskers of constant intensity. Each masker had a single 1/3-octave passband. The curve relating threshold to the central SF of the masking noise can potentially depend both on the tuning of SF mechanisms and on the occurrence of off-frequency viewing. For each signal SF, there were four maskers, centered at -1, -1/2, +1/2, and +1 octaves relative to signal SF. The noise intensity was again 37.3 dB *re* 0.1 min arc, so that by combining the present data with data from Experiment 1, one has data for five noise placements, the fifth being 0 octaves *re* signal.

Results and Discussion. Subject AG produced unstable, and sometimes unmeasurable, thresholds in some high-SF conditions. There were some problems for a 0.2 cpd signal masked by bandpass noise centered -1/2 oct *re* signal SF, but problems were particularly persistent for 1.1314 cpd and 1.6 cpd signals masked by bandpass noises centered +1/2 and +1 oct *re* signal. The problems seemed to arise during the practice trials, when incorrect responses produce large jumps (12 dB) in signal intensity. A few successive incorrect responses during this period can cause stimulus intensity to rise so high that depth is extremely poorly perceived, so that the discrimination cannot be performed on the basis of detecting a disparity grating. If the situation is not resolved before experimental trials begin, the subject may perform near chance level until all ten turnarounds occur. Nonmonotonic psychometric functions can result, and if they are forced into the Procrustean bed of threshold determination, the data-fitting procedure usually yields extremely high threshold estimates that can vary widely from run to run, and sometimes the data-fitting procedure fails to converge at all. Although AC did not produce such unstable thresholds, he did report difficulties a few times with extremely high intensity stimuli. Presumably, these difficulties all occurred during practice trials and were resolved before experimental trials began.

High-frequency conditions seem particularly prone to these problems--AC never reported difficulty in any but the highest-SF conditions. It is possible that this may be related to the disparity gradient (DG) limit described by Burt and Julesz (1980a, b), whereby a slanted stereoscopically presented surface is binocularly unfusible if it recedes in depth at a rate greater than one unit of disparity per unit of visual angle. Although the signal-plus-noise stimuli possess complex spectra, some insight to the problem can be gained by considering a sinusoidal disparity grating presented alone. For a grating of amplitude A and spatial frequency f, the maximum DG is $2\pi f A$. Because the DG limit is

1, a grating of spatial frequency f is just fusible when its amplitude is $A = 1/(2\pi f)$. For a 0.2 cpd grating, this amplitude is 47.7 min, or 50.6 dB *re* 0.1 min RMS; for a 1.6 cpd grating, this amplitude is only 6.0 min, or 32.5 dB *re* 0.1 min RMS. This latter value is in the range of threshold intensity for a grating masked by a noise of similar SF in the present experiments. Even the former, higher, intensity may be reached if the subject is distracted during the practice trials.

In order to obtain data from AG in all conditions, the conditions that yielded the unstable data were all re-run with the maskers' RMS amplitudes reduced by half (6 dB down). It is thresholds from these conditions, plus 6 dB to make them directly comparable to the other masking data, that are reported for AG for the unstable conditions.

The raw data were fit as in Experiment 2. The triangles in Figures 9-11 present the 75% thresholds thus derived. Threshold is generally highest where the masker is centered at the SF of the signal, and threshold generally decreases with increasing difference between signal SF and center SF of the noise. In addition, in almost all cases, the masking curves for bandpass noise are narrower than the corresponding curves for notched noise. For 24 out of 32 pairs of data points, the threshold in the bandpass noise condition falls below the threshold in the corresponding notched noise condition. In addition, the highest thresholds generally occur when the masker is centered at the SF of the signal. The exception is the highpass characteristic of AG's bandpass masking curve for a 1.6 cpd signal, which is discussed in more detail below, in Discussion. It is important to note, however, that there is not an indication that as signal SF is swept across the bank of filters, a difference between notched masking curves and bandpass masking curves appears, disappears, and reappears, as would be expected if there were a small number of channels. These features of the data support the conclusion that a continuum of SF channels

operates in detection of sinusoidal disparity gratings.

Another feature of the data concerns the symmetry of falloff. Maskers were shifted in 1/2-octave steps. It is unclear in Figures 9-11 if masking falls off more quickly in dB/octave as maskers shift toward lower frequencies *re* signal, or if masking falls off at the same rate in dB/octave whether maskers shift toward lower or higher frequencies *re* signal--the data are equivocal. However, it is clear that masking falls off at least as quickly as maskers shift toward lower frequencies as it does when maskers shift toward higher frequencies. This is incompatible with the psychophysical filters' being symmetrical on linear coordinates, because a constant-octave shift in frequency corresponds to a small change in linear frequency when the shift is toward lower frequencies, but it corresponds to a large shifts in linear frequency when the shift is toward higher frequencies.

Experiment 4. Absolute Threshold

Subjects. AC and DR served as subjects in all this experiment's conditions.

Design. In Experiment 4, thresholds were measured for the detection of sinusoidal disparity gratings in the absence of noise. In the absence of masker noise in the stimulus, the noise in the psychophysical filters are governed in this condition by the filter's intrinsic noise. Thus, in conjunction with the filter shapes derived from the masking experiments, above, the results of this experiment were intended to be used to estimate the intrinsic noise in the filters.

Results and Discussion. The raw data were fit as in Experiments 2-3. Figure 12 displays 75% thresholds thusly derived for all seven sinusoid SFs. Peak sensitivity is in the range of 0.28 to 0.4 cpd. This is in agreement with the results of Tyler and Julesz (1978), Schumer and Ganz (1979), and Rogers and Graham (1982). The high-frequency rolloff in sensitivity agrees with previous findings, though the specific shape differs. Given that the results of Schumer and Ganz (1979) and Rogers and Graham (1982) themselves differ

from one another in this respect (The results of Tyler & Julesz are unavailable in this detail), the lack of exact agreement, which could be due to inter-subject variability or specific characteristics of the stimuli (e.g., Schumer and Ganz used dynamic RDSs), is not troubling.

Analysis of Thresholds in Detection Experiments

The task of modeling the threshold data is equivalent to modeling the SF filtering that the stimuli undergo before they are submitted to the detector model developed in Appendix B and summarized in Experiment 1.

Based on the gross features of the data, two classes of models can be ruled out. A single-channel model, which asserts that the stereovisual system's plot of absolute threshold versus SF (CSF analog) describes the entirety of the system's SF selectivity, is clearly incompatible with the results of Experiments 2-4, because the data exhibit numerous masking curves (Experiments 2-3) much narrower than the overall system's CSF analog (Experiment 4). A model with a small number of channels also can be ruled out, based on the results of Experiments 2-3. Such a model predicts that the notched-noise masking curves would be just as narrow as the bandpass masking curves at some SFs, whereas the data from Experiments 2-3 indicate that bandpass masking curves are always narrower, with a sole exception in AG's data set, the significance of which is addressed in Discussion.

The remaining alternative is a model with a continuum of channels, and it is that alternative which is elaborated below. In developing the model, the data from Experiment 1 are used to determine the SNR at which a stimulus is at an observer's threshold. The data from Experiment 2-3 are used to determine the shape of the psychophysical filters, including how SF selectivity varies with SF. The data from Experiment 4 are used to determine the characteristics of the system's intrinsic noise.

Elaboration of the Quantitative Model. Assume that the observer's detection process can access any one of a family of psychophysical filters. Characterize this family with a set of real, non-negative functions in the SF domain, $\{|\phi_\nu(f)|^2\}$, that represent the squared transfer functions of filters with various center frequencies, ν. Normalize these functions so that their maximum value is $|\phi_\nu(\nu)|^2 = 1$. Because we consider only real valued input, all frequencies are non-negative. If there were a small number of filters, then ν could take on only a small number of values; however, given the evidence from Experiments 2 and 3 (and Experiment 5, below), there must be a continuum of filters, and ν can take on any real value in some interval. The energy detector incorporates filters in quadrature phase, and the output of any filtering operation is thus phase-insensitive. Assume that, before a stimulus undergoes analysis by an SF filter, it is degraded by noise intrinsic to the visual system. Given a stimulus with Fourier transform $X(f)$, the output of filter ν is

$$\Phi_\nu(f) = \int |\phi_\nu(f)|^2 |X(f)|^2 \, df + n_\nu, \tag{3}$$

where n_ν is a random variable representing the intrinsic noise, which is approximately distributed as $\sigma_\nu^2 \chi^2(2)$, where σ_ν^2 is a constant representing noise intensity, and $\chi^2(2)$ is the distribution of a chi-square random variable with 2 degrees of freedom. (Here, and subsequently, all integrations are from 0 to ∞ unless otherwise indicated.) Such a distribution is chosen because it represents the same approximation to the distribution of energy in a gaussian noise sample derived in Appendix B.

The detectability of a signal in noise is determined by the ratio of signal energy passed by the SF filter to the average noise energy passed the by SF filter (see Appendix B). The higher this *signal-to-noise ratio* (SNR), the more detectable is the signal.

In the notched-noise experiment (Experiment 2), best performance is achieved by

basing the detection decision on the output of a filter at the frequency, ς , of the signal of intensity S. For the sufficiently intense maskers used herein, the noise passed by the channel is dominated by the masker rather than by the channel's intrinsic noise, and the SNR is

$$S / \int |\phi_v(f)|^2 \Gamma_N(f) df, \tag{4}$$

where $\Gamma_N(f)$ represents the power spectrum of the masker. Because $\Gamma_N(f)$ is a known function determined by the SF characteristics of the stimulus, the notched noise data allow one to determine a $\{|\phi_v(f)|^2\}$ to model the shape of the psychophysical filters.

In the bandpass noise experiment (Experiment 3), the observer can exploit off-frequency viewing. In this experiment, a signal of intensity S and frequency ς is not optimally detected by basing detection on the output of filter $\phi_\varsigma(\cdot)$. Instead, the optimal filter is that whose squared transfer function, $|\phi_\xi(f)|^2$, maximizes the SNR,

$$\frac{|\phi_\xi(\varsigma)|^2 S}{\int |\phi_\xi(f)|^2 \Gamma_N(f) df}. \tag{5}$$

Model of Filter Shape. There are a multitude of functions with which one could model the shape of the psychophysical filters tuned for the SF of disparity modulation. The critical feature, in order to be able to account for the advantages of off-frequency viewing, is that the filter be flatter near its center frequency than it is on its skirts. A second feature is that the function be at least approximately symmetrical on logarithmic frequency coordinates. It is clear from the results of Experiment 3 that the filters are not symmetrical on linear frequency coordinates. As to the question of whether they are symmetrical or asymmetrical on logarithmic frequency coordinates, the data of Experiment 3 are equivocal. However, if a model were to capture any asymmetry, if ought to allocate

a free parameter to fit it, and there are too few data points at each signal SF to permit allocation of free parameters in such a spendthrift fashion. It is parsimonious in general to make simplifying assumptions unless they are clearly unwarranted. Thus, logarithmic filter symmetry will be assumed in the present model. A third feature, justifiable on grounds of parsimony and functional utility, is that the filter's falloff be monotonic. That is, the filter should possess no secondary passbands at frequencies distant from the center frequency. A disparity SF filter, as herein conceived, is selective for a narrow range of frequencies. Secondary passbands would defeat the filter's purpose, and positing them is unparsimonious in the absence of positive evidence for their existence. A fourth feature desirable of the function is that be smooth. Although ideal filters have infinitely steep and infinitely sharp skirts, it is unlikely that the human visual system can implement that ideal. Biologically realizable filters are probably smoother and have more gradual falloffs.

The model of filter shape chosen herein is

$$|\phi_{\nu}(f)|^2 = \begin{cases} \left(\dfrac{f}{2^{-p}\nu}\right)^{1/s}, & \text{for} f \le 2^{-p}\nu \\ 1, & \text{for } 2^{-p}\nu < f \le 2^{p}\nu, \\ \left(\dfrac{f}{2^{p}\nu}\right)^{-1/s}, & \text{for} f > 2^{p}\nu \end{cases} \tag{6}$$

where $s > 0$ is a parameter determining the slope of the falloff in the skirts, and $p \ge 0$ is a parameter determining the width of the flat portion (the passband). Specifically, the falloff in the skirts is approximately $3/s$ dB/oct, and the width of the flat portion is $2p$ oct. The particular parameterization of Eq. 6 in terms of s and p was chosen from among alternative, isomorphic parameterizations in order that the two parameters both have units

of octaves. The filter, whose profile resembles a trapezoid on logarithmic frequency coordinates, is depicted in broken lines in Figure 13. This function possesses the critical feature of being flat near its peak. It also possesses the desirable feature of logarithmic symmetry in frequency. In addition, it is also free of secondary lobes. The function, however, lacks the desirable feature of smoothness in its abrupt transition from passband to skirt. There are other functions that claim all the advantages of the trapezoidal function, but do not suffer the disadvantage of a nonsmooth transition. For example, a Butterworth filter, depicted in solid lines in Figure 14, has a smooth transition to its skirts, in spite of being very flat in the passband. The justification for choosing the trapezoid is twofold. For one, it is a reasonable approximation to other filters, such as the Butterworth filter, which is perhaps more biologically plausible. Second, it is mathematically and computationally tractable for the modeling enterprise herein. Because it is utterly flat within its passband, one can make the simplifying assumption in Experiment 1 that a 1/3-octave noise centered at the signal's SF is unattenuated, without the model being internally inconsistent. In addition, performing the integrations in Eqs. 4-5 is very easy with the trapezoidal function, whereas the corresponding integrations with a Butterworth filter cannot be found in closed form. Moreover, simulating off-frequency viewing by solving Eq. 5 for ξ is very easy with the trapezoidal function: given the assumption that the sharpness of filter tuning changes only gradually with SF, the best filter placement for detecting a signal in the presence of bandpass noise not centered at the signal's SF is one that maximally attenuates the noise without attenuating the signal. This filter is the one whose transition from passband to skirt falls at the signal frequency. More precisely, the best center frequency for detecting a signal of frequency ς in the presence of a narrowband noise centered at frequency $\zeta_{hi} > \varsigma$ is $\xi = 2^{-p}\varsigma$, and the best center frequency for detecting the same signal in the presence of a narrowband noise centered at

frequency $\zeta_{lo} < \varsigma$ is $\xi = 2^{p}\varsigma$. The advantage of the trapezoidal filter model is that these equations for ξ are very simple. Another choice of filter shape would require that, for each iteration of curve fitting, a nested cycle of iterations be performed in order to find that center frequency that maximizes the SNR.

For each signal SF, the threshold data from Experiments 2 and 3 were jointly fit, according to a least-squares criterion, with the predictions due to Eqs. 4-6, scaled upward downward individually for each subject to account for that subject's performance relative to the near-ideal observer's performance (see results of Experiment 1). Figures 14-16 and Figures 17-19 display redisplay the notched-noise and bandpass-noise data, respectively, along with the model's predictions. Tables 2-4 list the best-fitting parameters, the mean squared errors, and the 3 dB octave bandwidths of the corresponding filters. A filter's 3 dB bandwidth is its width at the point where the squared magnitude is half-maximum.

The 3 dB octave bandwidth of the trapezoid filter is

$$BW_{3\,dB} = 2p + 2s, \tag{7}$$

the first term of which is the octave bandwidth of the flat portion, and the second term of which is the 3 dB octave bandwidth of a filter with a skirt of the same steepness as the trapezoid filter, but with no flat portion. Such a filter would have a triangular profile on log-log coordinates.

Figure 20 depicts this triangular filter, along with the trapezoidal filter best suited for sinusoid detection in the presence of bandpass noise. The two filters differ from one another only where there is no stimulus energy. This section includes the entirety of the flat portion of the trapezoidal filter. Thus, the effect of off-frequency viewing is to remove the contribution of the flat portion of the filter to the falloff of masking, and the bandpass masking thresholds depend only on the slope of the skirt. Because the flat portion of the filter does not contribute to the falloff of bandpass masking, the bandpass data alone

cannot yield a measurement of the width of the flat portion of the filter. Analyzing the bandpass data, while neglecting the contributions off-frequency viewing, would therefore yield bandwidth estimates too low by the width of the flat portion. Such an analysis would thus yield only $2s$, the second term of Eq. 7, for the 3 dB octave bandwidth of psychophysical SF filters. The values Tables 2-4 list for $2s$ are indeed in rough agreement with Tyler's (1983) estimates of the bandwidth of SF-selective mechanisms, the narrowest of which was 1/3 to 1/2 octaves, compared with $2s$ as low as 0.34 oct herein. The values Tables 2-4 list for $2p$, the first term of Eq. 7, indicate the extent to which neglecting off-frequency viewing led to spuriously low bandwidth estimates.

Analysis of Absolute Detection Thresholds. According to the model of stimulus processing presented above, noise of approximate distribution $\sigma_v^2\chi^2(2)$ is added to the stimulus before an SF channel filters it. This represents degradation by noise intrinsic to the stereovisual system. In the masking experiments, this intrinsic noise was neglected, because its intensity was far lower than the intensity of the masking noise in the stimulus itself. However, in the absolute detection experiment (Experiment 4), there is no extrinsic masking noise. The intrinsic noise therefore limits detectability. In this task, assume best performance is achieved by basing the detection decision on the output of a filter at the frequency, ς , of the signal of intensity S. (In the model elaborated in the paragraph below, off-frequency viewing would only lower the SNR.) This filter leaves the signal unattenuated, but attenuates the noise. The signal-to-noise ratio here is

$$S/\sigma_v^2. \tag{8}$$

Before predicting absolute detection thresholds, additional data reduction was performed. In the analysis of Experiments 2-3, the parameters, s and p, were allowed to vary between signal SF conditions. Because it is reasonable to assume that the parameters change in a simple fashion with SF, the plots of s and p versus signal SF were fit with

polynomials. The variation in s was well fit by a second-order polynomial in $\log_2 \nu$, represented by the broken lines in Figure 21. The variation in p was also fit well by a second-order polynomial in $\log_2 \nu$, if the data at 1.6 cpd are neglected. We neglect these data on two grounds: 1) Both subjects who performed in high-SF conditions (AC and AG) occasionally reported difficulty: when intensity was too high, the stimulus was difficult to perceive. 2) This produced in some cases a nonmonotonic psychometric function (see Experiment 3, Results and Discussion). These effects, while worthy of investigation, probably reflect the action of mechanisms other than disparity SF filters (e.g., disparity-gradient limit, Burt & Julesz, 1980a, b).

The specific polynomials were

$$s = 0.434 + 0.184 \log_2 \nu + 0.048 \log_2^2 \nu$$

and (9)

$$p = 0.638 + 0.361 \log_2 \nu + 0.110 \log_2^2 \nu$$

for subject AC, and

$$s = 0.378 + 0.184 \log_2 \nu + 0.049 \log_2^2 \nu$$

and (10)

$$p = 0.638 + 0.859 \log_2 \nu + 0.353 \log_2^2 \nu$$

for subject AG. Thus, the 14 degrees of freedom (2 parameters × 7 SFs) were reduced to 6 (3 for each parameter), and simple functions were obtained for the two parameters. Tables 5-6 present AC's and AG's respective RMS errors for the original 14 df fit, along with the RMS errors for the 6 df fit. The overall RMS errors increase slightly, from 3.03 dB to 3.85 dB for AC, and from 5.16 dB to 6.48 dB for AG. An important consequence of this procedure is that it makes possible extrapolation of the parameter values beyond the range of conditions examined in the present paper; thus, it expands the opportunity for prediction of data yet to be collected. Henceforth, all calculations involving s and p will

use the values derived from these polynomials.

To predict the absolute detection thresholds, it is necessary to posit how σ_ν^2, the intensity of intrinsic noise, varies with ν , the center frequency of the filter. One alternative is that the intrinsic noise intensity varies with the filter's bandwidth in log frequency, so that narrowly tuned filters have little intrinsic noise and broadly tuned filters have more intrinsic noise. Finding a filter's bandwidth in log frequency is equivalent to finding how much noise of constant energy per octave it passes. Such a noise, which has spectrum level inversely proportional to frequency, is known as pink noise, because it includes a broad range of wavelengths, but has more energy in the low-frequency (long wavelength, "red") end of the spectrum. Thus, it is analogous to desaturated red light, i.e., pink light.

The noise intensity in the channel centered at frequency ν is therefore posited to be

$$\sigma_\nu^2 = \left(\int |\phi_\nu(f)|^2 N_1 f^{-1} \, df \right)^\varepsilon , \tag{11}$$

where N_1 is a free parameter determining the pink noise's energy per octave, and ε is a free parameter determining how rapidly σ_ν^2 varies with the energy of pink noise passed. Evaluating Eq. 11 for the trapezoidal filter defined in Eq. 6 and substituting the result into the SNR equation (Eq. 8) yields an SNR of

$$S \left(2s + 2p \ln 2\right)^{-\varepsilon} N_1^{-\varepsilon} , \tag{12}$$

the predictions of which are plotted in Figure 12, along with the data. Because subject DR did not participate in most of Experiments 2-3, from which are derived the polynomial approximations to the trapezoidal filter's parameters, subject AC's polynomials (Eq. 9) were used to derive predictions for both subjects. Table 7 presents the best-fitting parameter values and corresponding RMS errors. The fit is good, with a root mean

squared error of 0.89 dB, pooled across the two subjects. Minimum threshold is predicted to occur at the center frequency of the filter that passes the least pink noise. According to the polynomial approximations in Eq. 9, this frequency is at 0.296 cpd.

Summary of Model of SF Selectivity. To summarize, a model incorporating a continuum of filters was developed with the data from the detection experiments. The threshold-versus-masking-intensity results (Experiment 1) were consistent with the existence of linear channels, and they showed that human observers' threshold SNRs were near ideal, and that information was pooled over no more than 2.7 cycles. The masking curves in the notched noise (Experiment 2) and bandpass noise (Experiment 3) experiments were used to generate a trapezoidal model of the psychophysical filter. Detection was assumed to be based on a filter tuned to an SF that maximizes the SNR. In the notched noise experiment, this filter is centered at the signal's SF. In the bandpass noise experiment, this filter is centered on the side of the signal's SF opposite from the masker's center SF. The model thusly quantifies the phenomenon known as off-frequency viewing. The characteristics of the trapezoidal filter that allow it to successfully predict the extent off-frequency viewing are its flat passband and steeper skirts. The filter's parameter that describes the slope of the skirts was found to vary in a parabolic fashion with center SF; the parameter that describes the width of the passband was found to be roughly constant with changing center SF. The trapezoidal filter accounted for the absolute threshold data (Experiment 4) well, if combined with the assumption that narrowly tuned filters had less intrinsic noise than broadly tuned filters.

Experiment 5. SF Discrimination

In this last experiment, thresholds were measured for the discrimination of sinusoidal disparity gratings of equal intensity according to their difference in SF. This shape-discrimination task differs in kind from the previous four experiments, all of which

assessed grating detection.

Subjects

AC served as subject in this experiment.

Design

The SF discrimination experiment proceeded as the detection experiments did, with the same apparatus, viewing conditions, and psychophysical procedures, with the following exceptions:

In each trial, the subject's task was to indicate in which of two intervals the sinusoid of higher SF was presented. Within a block of trials, the geometric mean of the sinusoid SFs was held constant. Mean SF ranged from 0.2 cpd to 1.6 cpd, in half-octave steps. These match the SFs of the signals in the detection experiments.

Sinusoid intensity was held constant within a trial block, during each of which two independent adaptive staircases were randomly interleaved. The signals were presented at 37 dB or 43 dB *re* 0.1 arcmin. The subject completed two blocks of trials in each of his 14 stimulus conditions (7 SFs $\times$ 2 intensities).

During a block of trials, the relative difference in spatial frequency, $\Delta f/f$, was varied in logarithmic steps. Each staircase's initial value for $\ln(\Delta f/f)$ was -3, corresponding to $\Delta f / f \approx 5\%$. During a staircase's practice trials, which continued until six turnarounds in that staircase were completed, $\ln(\Delta f/f)$ was decreased by 0.25 following correct responses and increased by 1 following incorrect responses. Following the practice trials, $\ln(\Delta f/f)$ was decreased by 0.125 following correct responses and increased by 0.5 following incorrect responses. Each staircase continued after practice until ten additional turnarounds in that staircase were completed.

Results

Threshold estimates were taken as the mean of the last ten turnarounds in each

staircase. These thresholds are displayed in Figures 23. It is clear from the upper panel in Figure 23 that there was no difference in SF discrimination between the 37 dB and 43 dB intensities. Performance at 37 dB exceeded that at 43 dB as often as performance at 43 dB exceed that at 37 dB. For this reason, results were pooled across intensity.

Analysis

If SF channels underlie SF discrimination, then the filter model ought to predict these results. This is the same approach used by Wilson and Gelb (1984), who modeled the discrimination of the SF of luminance modulation with a six-channel line-element model (Wilson, McFarlane, & Phillips, 1983). A model with a discrete number of channels predicts that SF discriminability will be best at SFs where a channel has steep skirts, because it is at these locations where a small change in SF results in a large change in channel activation. Where there are no steep skirts, SF discriminability ought to be poor. Thus, such a model predicts that SF difference limen varies with SF in a scalloped fashion, that a plot of SF difference limen versus SF ought to alternate between convex and concave regions. No such alternation is evident in the results in Figure 23. This lends additional support to the rejection of the hypothesis that there is a small number of SF channels.

The results from the absolute detection experiment (Experiment 4) provide estimates of the intrinsic noise in the filters, as a function of filter's center frequency. If this noise were the limiting factor in SF discrimination, then the difference limen for SF would decrease with increasing signal intensity. This is because the higher a signal's intensity, the higher is the SNR after that signal is filtered. If SF discrimination is based on the output of SF filters, then high SNRs enable precise discrimination. However, no differences were found in SF discriminability between 43 dB signals and 37 dB signals, even though SNR varies by 6 dB between these conditions. Thus, intrinsic noise in SF

channels cannot be the limiting factor in SF discrimination.

Because some factor unrevealed by the detection experiments must be limiting SF discrimination, it would be problematic to motivate a model of SF discrimination based on signal-to-noise ratios. Instead, we follow the lead of Wilson and Gelb (1984), who forgo SNR arguments in favor of relating proportion correct in SF discrimination directly to differences in a filter's response to the two stimuli of different spatial frequency.

Difference limen for SF was assumed to be based on the output of a filter tuned so that the sinusoid SFs fell near the top of a filter's skirt, close to the filter's elbow. For a given average sinusoid SF, there are two filters so tuned: one whose center frequency is above sinusoid SF and another whose center frequency is below sinusoid SF. Discrimination was assumed to be based on the output of the lower-SF filter, although results are similar if discrimination is assumed to be based on the output of the higher-SF filter or based on the output of whichever filter has the steeper skirt. As derived in Appendix C, threshold is

$$\ln\left(\left.\frac{\Delta f}{f}\right|_{\text{thresh}}\right) = \ln\left[-1+2^{\left(s^{\eta}/a\right)}\right], \tag{13}$$

where a and η are free parameters, and s is the skirt parameter of the trapezoidal filter on whose output SF discrimination is based. The pooled data were fit with Eq. 13. The dashed line in Figure 23's right panel displays the predictions. The best-fitting parameter values were a=0.84 and $\eta = 2.15$. These yielded an RMS error of 0.196, accounting for 79% of the variance in the pooled values of $\ln(\Delta f/f)$, and corresponding to a 22% relative error in $\Delta f/f$ values. Minimum difference limen was predicted to fall at 0.314 cpd.

Discussion

The goal of this research was to investigate the selectivity of the stereovisual

system for the spatial frequency of binocular disparity modulation, with special attention to the questions: What are the channels' bandwidths?; What is the extent of off-frequency viewing, and does it disappear at some SFs?; Is there a small number of channels or is there a continuum?; and Is the channel idea applicable across experiments in different domains? The approach taken in this research was conduct detection experiments and use the results to develop a quantitative model of the SF selectivity.

The difference between notched masking curves and bandpass masking curves indicated that there was a continuum of channels with flat passbands and disconfirmed the hypothesis that there was a small number of channels. The small-number-of-channel hypothesis was further disconfirmed by the lack of scalloping in Experiment 5's SF discrimination plot (Figure 23). Because such a continuum of channels admits off-frequency viewing at every SF, bandwidth estimates from previous experiments that used bandpass pseudo-noise maskers (Tyler, 1983) were too narrow, because they did not take into account the way off-frequency viewing yields masking curves narrower than the channels' tuning curves. The model quantifies the extent to which off-frequency viewing narrows bandpass masking curves. According to Eqs. 9-10's polynomial estimates of the passband parameter, p, which are displayed in Figure 22, off-frequency viewing narrows bandpass masking curves by at least about a quarter octave (For subject AC, $2p = 0.68$ oct at 0.32 cpd; for subject AG, $2p = 0.23$ oct at 0.43 cpd). The extremely narrow tuning as low as 1/3 to 1/2 that has been reported in a previous study (Tyler, 1983) is roughly consistent with the smallest values for $2s$ that Eqs. 10-11 (Figure 21) yield: For subject AC, $2s = 0.26$ oct at 0.26 cpd, and for subject AG, $2s = 0.21$ oct at 0.27 cpd. In the context of the model, this "tuning" would be regarded as spuriously narrow due to the effect of off-frequency viewing. When off-frequency viewing is taken into account, channels are found to be tuned no more sharply than 1.2 oct (For subject

AC, $2p+2s=1.20$ oct at 0.30 cpd; for subject AC, $2p+2s=1.41$ oct at 0.41 cpd).

However, in light of the present results on the operation of off-frequency viewing, the unpublished results of Tyler that masking curves yielded by notched pseudo-noises are just as wide as masking curves yielded by bandpass pseudo-noises is puzzling. Evaluation of the apparent discrepancy will require obtaining the data and exact stimulus conditions and computing the predictions of the model.

The question about the extent to which the channel idea is applicable to different domains was addressed by successfully applying the model to the results of a 3D shape-discrimination experiment, in which SF discrimination was measured. Although this represents but one of a multitude of domains, it nevertheless bodes well for the theoretical power that the disparity-SF channel idea will ultimately demonstrate.

High SFs

As was reported in Experiment 3-4, Results and Discussion, subjects sometimes experienced difficulty in high-SF conditions. Although AG's persistent difficulties at the highest SFs were alleviated somewhat by lowering masker intensity, the highest SF maskers still generated very high thresholds, to the point where AG's masking curve for 1.6 cpd signal in bandpass noise appears highpass, not bandpass. Consistent with these findings, Tyler's (1983) results show that even for signals of only 0.6 cpd, the masking curve does not fall very quickly for high-SF maskers. Moreover, when the masker was reduced in amplitude by a factor of three (9.54 dB down), the masking curve had a steeper high-SF rolloff than it had before the masker amplitude was reduced (Tyler, 1983). As was discussed above, in Experiment 3-4, Results and Discussion, the results may be due to fusion problems related to the DG limit (Burt & Julesz, 1980a, b). Another possibility also presents itself. Whereas in most linear systems, one would expect frequencies beyond a systems range to either be attenuated or to produce aliasing, high-SF disparity

modulations of sizable intensity do neither. Instead, transparency results: When the grating is not resolvable, dots with crossed disparities and dots with uncrossed are simultaneously perceived. The peak and trough disparities are resolvable, but the disparity modulation is not. A fundamental assumption of the present analysis is that disparity can be represented as a scalar field. When more than one disparity is perceived at the same location, that assumption is violated. This would define a limit on the range of validity of the SF channel description. To investigate this explanation would require a systematic investigation of when the transparency phenomenon emerges as intensity and SF characteristics of a disparity modulation change.

Low SFs

In addition to studying high SFs, lower SFs warrant further study. The present results indicated that filters were most narrowly tuned about 0.2 to 0.4 cpd. That is very close to the lowest SF tested. That makes the estimate of narrowest tuning unstable, because data are largely missing on the low-SF side. The issue is particularly important in order to resolve how the width of the passband and the slope of the skirts covary. Both parameters were found to vary parabolically with logarithmic frequency. It is parsimonious to assume that both parameters reach a minimum at the same SF, so that filters with narrow passbands have steep skirts. However, without low-SF data, estimates of the SF at which these parameters are minimized are unstable.

Spatial Extent

In the present study, spatial extent of the stimulus was not systematically manipulated. It was instead held constant, with the consequence that in higher-SF conditions, more cycles of the stimulus were visible than in lower-SF conditions. If SF mechanisms pooled information over a large spatial extent, this would cause thresholds to be lower in the higher-SF conditions. However, when signals were presented in the

presence of narrowband noises centered at the signal's SF, threshold was found to be independent of SF (Experiment 1). This indicates that the extent of spatial pooling is upper bounded at the 2.7 cycles visible in the 0.2 cpd signal condition. Further studies could manipulate number of cycles directly in order to assess the range of spatial pooling.

In-Depth Study of Filter Shape

A number of simplifying assumptions were made about the filter shape in the present model. For one, the filter was assumed to be perfectly flat in the passband. For another, it was assumed to have a uniform exponential rolloff at the same rate on the high-SF and low-SF skirt, and the skirts were assumed to join sharply to the passband. In addition to being made for computational and analytical reasons, these assumptions were made because there were relatively few masker SFs tested for each signal SF, and thus relatively few degrees of freedom at any given signal frequency. The study was designed this way because, in order to test the hypothesis that there was a small number of channels, it was important to spend subject time collecting data at several SFs. With that question resolved, another study could be undertaken that examines but a few signal SFs, but uses many different maskers and collects enough data to yield highly reliable estimates of the shape of the filter. This was the approach taken by Patterson and Nimmo-Smith (1980) who, because they were investigating a psychoacoustic filter, did not have to concern themselves with the question of whether there was only a small number of filters. An experiment such as this could yield reliable information on such questions as whether the filter symmetrical in logarithmic SF (The results of Experiment 3 already eliminate symmetry in linear SF).

Relation of Disparity SF Channels to Motion Parallax.

It was noted in the Introduction that disparity and motion parallax were both cues for object recognition. In addition to this functional link between the cues, there is also a

close formal similarity in the cues: Assessing binocular disparity is a matter of detecting small spatial displacements in a corresponding part of an object's simultaneous projection to the two eyes; assessing motion parallax is a matter of detecting small spatial displacements in a corresponding part of an object's successive projection to either eye. Finally, in previous research, it has been proposed that the two cues are linked in mechanism. An early suggestive finding was that, when measured with similar apparatus and psychophysical procedures, the curve describing threshold for disparity modulation as a function of SF was very similar to the curve describing threshold for parallax modulation as a function of SF (Rogers & Graham, 1982). More recently, sub-threshold interactions have been found between disparity gratings and parallax gratings (Bradshaw & Rogers, 1992). To further evaluate the relationship between the processing of binocular disparity and motion parallax, the present study's experiments and model can be directly adapted for motion parallax by simply substituting amplitude of parallax modulation for amplitude of disparity modulation. If the present study's pattern of results were duplicated in the parallax study, it would be consistent with similar mechanisms or shared mechanisms. A further test of the extent to which the two cues are processed by the same mechanisms would be to perform a cross-masking experiment, whereby one would measure threshold for detection of a disparity grating in the presence of parallax noise, or *vice versa.* If the masking curve was the same shape as the masking curves for disparity gratings masked by disparity noises and for parallax gratings masked by parallax noises, then there would be strong evidence for a common mechanism in the processing of the two cues.

Function of Disparity SF Channels.

Because disparity SF channels detect patterns of depth variation, they are well suited as an early stage in 3D object recognition. But, what is their more specific function? That is, *how* do disparity SF channels contribute to object recognition?

Answering this functional question is more speculative than answering such structural questions as "What is the BW?" However, it need not yield untestable propositions. The functional hypotheses proposed below can yield predictions to confront present and future data.

Multiple Spatial Scales. A set of channels tuned to different SFs represents depth variations across different spatial scales. By filtering an object's disparity profile with a low-SF channel, the visual system can extract its overall (coarse) shape. By filtering an object's depth profile with a high-SF channel, the visual system can extract fine detail of its shape.* The question arises, however, if these operations can be performed simultaneously on the same image. If so, can the visual system make use of this simultaneously available information? Answering this question would require the use a non-sinusoidal signal, so that the signal's energy is parceled between more than one filter. The results of Schumer and Ganz's (1979) sub-threshold summation experiment, wherein the detection of a two-component compound grating was determined by the independent detectabilities of the two components, indicates that the visual system can use the information in at least two channels simultaneously. Additional experiments, using signals with more complex spectra, are necessary to find out how many channels the visual system can simultaneously monitor.

Dimensional Reduction. The number of possible inputs to the stereovisual system

---

* The terms *low-SF*, *high-SF*, *coarse*, and *fine* are referred relative to the peak SF for sensitivity to disparity modulation. Because the SF of peak sensitivity for disparity modulation is so low (Schumer & Ganz, 1979; Rogers & Graham, 1982; Experiment 5), all visible disparity SFs are low or medium, relative to the SF of peak sensitivity for luminance modulation.

is enormous, and the number of dimensions required to describe all the possible inputs is likewise enormous. However, by having a few subsystems code stimulus characteristics on specific stimulus dimensions, the stereovisual system can perform dimensional reduction. This dimensional reduction would serve to extract only the important information from an image, most of whose information is unimportant. From this vantage point, the set of SF channels can be seen as a subsystem coding the stimulus characteristic of spatial frequency. Does this coding capture a large amount of information in the image? That is, is SF coding efficient? If so, then this coding performs data compression by reducing the redundancy in the encoding of the stereovisual information (Attneave, 1954; Barlow, 1961). The answer to the question depends on the characteristics of the image. If the image has little periodic structure, then SF coding would capture little information; if the image has much periodic structure, then SF coding would capture much information. Presumably, SF coding is efficient; otherwise, it would not have evolved. However, an empirical answer to the question would require a survey of depth patterns in the visual environment. There is reason to believe, however, that depth periodicities are indeed common, because periodicities in the structure of biological objects are common. These biological periodicities probably arise largely because it is efficient to genetically code how to build a large structure by coding how to build the repeating units that make it up. This presents the interesting likelihood that the efficiency of coding an image in terms of its periodicities arises ultimately from the efficiency of genetically coding biological structures themselves in terms of periodically appearing units.

Disparity Curvature. Disparity SF channels may function in computing the second spatial derivative of disparity, called disparity curvature. It is equivalent to the rate of change of disparity gradient (Burt & Julesz, 1980a, b) across space. The notion that humans can perceive disparity curvature has previously been proposed as a way to avoid

directly re-scaling perceived depth as the vergence angle changes, because, unlike binocular disparity, which varies approximately inversely with the square of viewing distance, disparity curvature is approximately invariant with viewing distance (Rogers, 1988).

Four lines of evidence indicate that depth curvature (the second spatial derivative of depth, whatever be the depth cue) is relevant for depth perception. First, when presented with figures whose depth information conveyed by binocular disparity and by monocular cues conflicts, perception is governed by a cue that depicts a curved surface over a cue that depicts a flat surface (Stevens & Brookes, 1988). Second, post-fusional latency for correct surface perception in random-dot stereograms (RDSs) is longer for flat surfaces than for surfaces whose depth curvature does not vanish everywhere (Gillam et al., 1988). Third, stereoacuity and depth judgments are degraded when the target stimuli are displayed embedded in, or in the presence of, representations of a flat object (Mitchison & Westheimer, 1984); performance in these tasks can be accounted for by assuming that the visual system is computing an approximation to depth curvature (Mitchison & Westheimer, 1990). Fourth, humans can accurately match the disparity curvature of two parabolic cylinders at different viewing distances, and the Weber fraction for disparity curvature is roughly constant over a two-and-a-half orders of magnitude range of disparity curvature (Rogers, & Cagenello, 1989).

The disparity curvature of a field can be approximated by operating upon the field with a linear operator with an excitatory main lobe in the middle and an inhibitory secondary lobe at each side. Because the Fourier transform of such a waveform has a bandpass characteristic, an SF channel could serve to approximate disparity curvature. To view the situation from a different perspective, a model that computes disparity curvature at different scales may be approximating the stereovisual system's SF analysis. One way to

test this hypothesis would be to determine if the model of SF selectivity developed herein can predict how the Weber fraction for disparity curvature varies with disparity curvature (Rogers & Cagenello, 1989). There are two reasons for optimism regarding this venture. The first reason is that disparity-curvature discrimination is a shape-discrimination task; the SF-discrimination task in Experiment 5 also represents a shape-discrimination task, and performance in that task was successfully predicted via the model of SF channels. The second reason is that the data of Rogers and Cagenello do appear to show some systematic variation in the Weber fraction for disparity curvature as disparity curvature changes. This suggests that disparity curvature *per se* may not be the relevant parameter. Because parabolic cylinders are not pure sinusoids, prediction of the Rogers-Cagenello data would require that the model be elaborated to describe how information is pooled across SF channels for the purposes of shape discrimination.

Laplacian. Another possible function for disparity SF channels derives from that fact that any finite resolution, the nonoriented second spatial derivative (Laplacian) of a field can be approximated by operating upon the field with a center-surround linear operator. Because the Fourier transforms of such waveforms have a bandpass characteristic, different SF channels may function to approximate the Laplacian of the disparity field at different spatial scales (cf., Marr, 1980). The utility of computing the Laplacian would be that by localizing its zero crossings, the visual system could perform 3D edge detection (cf., Marr & Hildreth, 1980; Marr, 1980). The Laplacian is an isotropic operator. That is, it is circularly symmetrical, and thus non-oriented. All studies of disparity SF have used stimuli with only a single orientation of periodic energy. Thus far, there has been a dearth of investigations of orientation tuning of hypercyclopean channels (for an exception, see Tyler's [1975] demonstration of a tilt aftereffect). To the author's knowledge, no study has assessed a channel's joint tuning for SF and orientation.

Just such a study would be required to determine if SF channels have the appropriate structure to approximate the Laplacian of the depth field.

In additional to being non-oriented, disparity SF channels would have to be spatially localized and repeating to approximate the Laplacian of a depth image. This is because the Laplacian of an image is not a single value. Instead, the Laplacian of an (x, y) image is itself a function of x and y coordinates. The results of Experiment 1 did indicate that the channels did not pool information over more than 2.7 cycles. Thus, they must be localized. However, no one has conducted an experiment to answer the question of whether disparity SF channels repeat across space. One experiment to answer the question would involve selectively adapting different locations in the visual field to different SFs by presenting grating patches (e.g., two-dimensional Gabor functions). If threshold for subsequently presented grating patches were jointly dependent on location and SF, relative to the adapting patches, then there would be evidence for repeating localized SF disparity channels.

Anatomical Blueprint. A final hypothesis about disparity SF channels is that they have no functional motivation at all. The idea is that center-surround antagonistic relations are common in the mechanism of the visual system, probably for functional reasons in many instances. This structure might be so widely applicable that the visual system evolved ~~to~~ reduplicate it on a large scale, with the result that some sub-systems to which the basic structure was neither particularly beneficial nor overwhelmingly detrimental would develop with the structure in place. Although this is a hypothesis of last resort, it would gain support if physiological, anatomical, and psychophysical evidence proliferated that the center-surround structure operated even in visual functions for which no purpose for the structure were reasonable.

Summary

This research involved a succession of four detection experiments, the results of which were used to develop a quantitative model of the visual system's selectivity for the SF of binocular disparity modulation. The model incorporated a continuum of linear filters with flat passbands, and a near-ideal detector that based decisions on the filter with the highest signal-to-noise ratio. Even though the number of cycles visible covaried with SF, when a signal was presented in narrowband noise centered at the signal's SF, threshold was found to be independent of SF. This indicates that disparity SF mechanisms do not pool over the more than 2.7 cycles visible in the 0.2 cpd condition.

If there were a small number of SF channels, then differences between notched-noise masking curves and bandpass masking curves would have successively appeared and disappeared as signal SF was swept across the filter bank. Because that result did not obtain, the model was based on a continuum of filters.

The model provided a quantitative estimate of the bandwidth of the channels underlying SF selectivity and of the effect of off-frequency viewing. Channels had minimum bandwidths of 1.20 oct, and off-frequency viewing narrowed bandpass masking curves by at least 1/4 octave.

The model was successfully applied to the results of a shape-discrimination experiment, in which SF discrimination was measured across SF. Hypotheses about how the SF channels function in 3D object recognition were proposed, and predictions and tests of these hypotheses were generated.

References

Attneave, F. (1954). Informational aspects of visual perception. *Psychological Review, 61*, 183-193.

Barlow, H. B. (1961). Three points about lateral inhibition. In W. A. Rosenblith (Ed.), *Sensory communication* (pp. 782-786). New York: Wiley.

Blake, R., & Holopigian, K. (1985). Orientation selectivity in cats and humans assessed by masking. *Vision Research, 25*, 1459-1467.

Blakemore, C., & Campbell, F. (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology, 203*, 237-260.

Blakemore, C. B., & Julesz, B. (1971). Stereoscopic depth aftereffect produced without monocular cues. *Science, 171*, 286-288.

Blakemore, C. B., & Sutton, P. (1969). Size adaptation: A new aftereffect. *Science, 166*, 245-247.

Braddick, O., Campbell, F. W., & Atkinson, J. (1978). Channels in vision: basic aspects. *Handbook of Sensory Physiology, 8*, 3-38.

Bradshaw, M. F., & Rogers, B. J. (1992). Subthreshold interactions between binocular disparity and motion parallax. *Investigative Ophthalmology and Visual Science, 33*(4), 1332.

Burt, P., & Julesz, B. (1980a). A disparity gradient limit for binocular fusion. *Science, 208*, 615-617.

Burt, P., & Julesz, B. (1980b). Modifications of the classical notion of Panum's fusional area. *Perception*, 671-682.

Dealy, R. S., & Tolhurst, D. J. (1974). Is spatial adaptation an aftereffect of prolonged inhibition? *Journal of Physiology, 241*, 261-270.

DeValois, R. L., Albrecht, D. G., & Thorell, L. G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research, 22*, 545-559.

DeValois, R. L., & DeValois, K. K. (1988). *Spatial vision.* New York: Oxford University Press.

DeValois, R. L., DeValois, K. K., Ready, J., & von Blanckensee, H. (1975). Spatial frequency tuning of macaque striate cortex cells. *Association for Research in Vision and Ophthalmology, 15*, 16 (A).

DeValois, R. L., Yund, E. W., & Hepler, N. (1982). The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research, 22*, 531-544.

Felleman, D. J., & Van Essen, D. C. (1987). Receptive field properties of neurons in area V3 of macaque monkey extrastriate cortex. *Journal of Neurophysiology, 57*, 889-920.

Gillam, B., Chambers, D., & Russo, T. (1988). Postfusional latency in stereoscopic slant perception and the primitives of stereopsis. *Journal of Experimental Psychology: Human Perception and Performance, 14*, 163-175.

Hubel, D H., & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *Journal of Physiology, 148*, 574-591.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology, 160*, 106-154.

Hubel, D. H., & Wiesel, T. N. (1970). Cells sensitive to binocular depth in area 18 of the macaque monkey cortex. *Nature, 225*, 41-42.

Johnson, N. L., & Kotz, S. (1970). Continuous univariate distributions. Boston: Houghton-Mifflin.

Julesz, B. (1960). Binocular depth perception of computer-generated patterns. *Bell*

System Technical Journal, *39*, 1125-1162.

Legge, G. E., & Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America*, *70*, 1458-1470.

Lehky, S. R., & Sejnowski, T. J. (1990). Neural model of stereoacuity and depth interpolation based on a distributed representation of stereo disparity. *Journal of Neuroscience*, *10*, 2281-2299.

Marr, D. (1982). *Vision*. New York: Freeman.

Marr, D., & Hildreth, E. (1980). Theory of edge detection. *Proceedings of the Royal Society of London B*, *207*, 187-217.

Mitchison, G. J., & Westheimer, G. (1984). The perception of depth in simple figures. *Vision Research*, *24*, 1063-1073.

Mitchison, G. J., & Westheimer, G. (1990). Viewing geometry and gradients of horizontal disparity. In C. Blakemore (Ed.), *Vision: coding and efficiency* (pp. 302-309). Cambridge, UK: Cambridge University Press.

Nachmias, J. (1967). Effect of exposure duration on visual contrast sensitivity with square-wave gratings. *Journal of the Optical Society of America*, *57*, 421-427.

O'Loughlin, B. J., & Moore, B. C. J. (1981). Off-frequency listening: Effects on psychoacoustical tuning curves obtained in simultaneous and forward masking. *Journal of the Acoustical Society of America*, *69*, 1119-1125.

Pantle, A., & Sekuler, R. (1968). Size-detecting mechanisms in human vision. *Science*, *162*, 1146-1148.

Patterson, R. D., & Nimmo-Smith, I. (1980). Off-frequency listening and auditory-filter symmetry. *Journal of the Acoustical Society of America*, *67*, 229-245.

Patnaik, P. B. (1949). The noncentral χ^2 and *F*- distributions and their applications. *Biometrika*, *36*, 202-232.

Poggio, G. F. (1984). Processing of stereoscopic information in primate visual cortex. In G. M. Edelman, W. E. Gall, & W. M. Cowan (Eds.), *Dynamic aspects of neocortical function.* (pp. 613-635) New York: Wiley.

Poggio, G. F., & Fisher, B. (1977). binocular interaction and depth sensitivity of striate and prestriate cortical neurons in the behaving rhesus monkey. *Journal of Neurophysiology, 40*, 1392-1405.

Poggio, G. F., Gonzalez, F., & Krause, F. (1988). Stereoscopic mechanisms in monkey visual cortex: binocular correlation and disparity selectivity. *Journal of Neuroscience, 8*, 4531-4550.

Poggio, G. F., & Talbot, W. H. (1981). Mechanisms of static and dynamic stereopsis in foveal cortex of the rhesus monkey. *Journal of Physiology, 315*, 469-492.

Quick, R. F. (1974). A vector-magnitude model of contrast detection. *Kybernetik, 16*, 65-67.

Rogers, B. (1988). Perspectives on movement. *Nature, 333*, 16-17.

Rogers, B., & Cagenello, R. (1989). Disparity curvature and the perception of three-dimensional surfaces. *Nature, 339*, 135-137.

Rogers, B., & Graham, M. (1982). Similarities between motion parallax and stereopsis in human depth perception. *Vision Research, 22*, 261-270.

Schiller, P. H., Finlay, B. L., & Volman (1976a). Quantitative studies of single-cell properties in monkey striate cortex. II. Orientation specificity and ocular dominance. *Journal of Neurophysiology, 39*, 1320-1333.

Schiller, P. H., Finlay, B. L., & Volman (1976b). Quantitative studies of single-cell properties in monkey striate cortex. III. Spatial frequency. *Journal of Neurophysiology, 39*, 1334-1351.

Schumer, R. A., & Ganz, L. (1979). Independent stereoscopic channels for different

extents of spatial pooling. *Vision Research, 19*, 1303-1314.

Stevens, K. A., & Brookes, A. (1988). Integrating stereopsis with monocular interpretations of planar surfaces. *Vision Research, 28*, 371-386.

Thomas, J. P., & Olzak, L. A. (1990). Cue summation in spatial discriminations. *Vision Research, 30*, 1865-1875.

Tyler, C. W. (1975). Stereoscopic tilt and size aftereffects. *Perception, 4*, 187-192.

Tyler, C. W., & Julesz, B. (1978). Spatial frequency tuning for disparity gratings in the cyclopean retina. *Journal of the Optical Society of America, 68*, 1365 (A).

Tyler, C. W. (1980). Functional analysis of cyclopean stereopsis. *Journal of the Optical Society of America, 70*, 1584 (A).

Tyler, C. W. (1983). Sensory processing of binocular disparity. In C. M. Schor & K. J. Ciuffreda (Eds.), *Vergence eye movements: Basic and clinical aspects*, Boston: Butterworth.

Wilson, H. R. (1980). A transducer function for threshold and supra-threshold human vision. *Biological Cybernetics, 38*, 171-178.

Wilson, H. R., & Bergen, J. R. (1979). A four mechanism model for threshold spatial vision. *Vision Research, 19*, 19-32.

Wilson, H. R., & Gelb, D. J. (1984). Modified line-element theory for spatial-frequency and width discrimination. *Journal of the Optical Society of America A, 1*, 124-131.

Wilson, H. R., McFarlane, D. K., & Phillips, G. C. (1983). Spatial tuning of orientation selective units estimated by oblique masking. *Vision Research, 23*, 873-882.

Appendix A

Autocorrelation Function of Narrowband Gaussian Noise

We wish to find the autocorrelation function, $\rho(y)$, of narrowband gaussian noise logarithmically centered about frequency f_g. Let the upper and lower cutoff frequencies of the noise be $c f_g$ and f_g/c, respectively ($c > 1$), and let the spectrum level of the noise be N_0. The arithmetic mean of the cutoff frequencies is thus $f_a = f_g(c+1/c)/2$, and the half-bandwidth is $f_{BW} = f_g(c-1/c)/2$. Denote the autocovariance function by $\gamma(t)$.

The Fourier transform of $\gamma(t)$ is the power spectrum,

$$\Gamma(f) = \tilde{\Gamma}(f) * [\delta(f - f_a) + \delta(f + f_a)]/2, \tag{A1}$$

where $*$ denotes convolution, $\delta(\cdot)$ represents the Dirac delta distribution, and

$$\tilde{\Gamma}(f) = \begin{cases} N_0, & \text{if } |f| < f_{BW} \\ 0, & \text{otherwise} \end{cases} \tag{A2}$$

is a rectangular pulse in frequency, with inverse Fourier transform $\tilde{\gamma}(t)$. $\Gamma(f)$ and $\tilde{\Gamma}(f)$ are plotted in Figures 24 and 25, respectively.

By the convolution theorem, the autocovariance function is thus

$$\begin{aligned} \gamma(y) &= \tilde{\gamma}(y)\mathbf{F}^{-1}\{[\delta(f - f_a) + \delta(f + f_a)]/2\} \\ &= \tilde{\gamma}(y)\cos(2\pi f_a y) \\ &= \frac{N_0}{\pi y}\sin(2\pi f_{BW} y)\cos(2t f_a y) \end{aligned} , \tag{A3}$$

where $\mathbf{F}^{-1}\{\cdot\}$ denotes the inverse Fourier transform, and the autocorrelation function is

$$\begin{aligned} \rho(y) &= \gamma(y)/\gamma(0) \\ &= \text{sinc}(2 f_{BW} y)\cos(2\pi f_a y)' \end{aligned} \tag{A4}$$

where the $\text{sinc}(\cdot)$ function is defined by

$$\mathrm{sinc}(q) = \frac{\sin(\pi q)}{\pi q}. \tag{A5}$$

The autocorrelation function, plotted in Figure 26, is thus a sinc-windowed cosine. Outside the central hump of the sinc window, the correlation is very low. We can thus define the *coherence length,* the distance beyond which instantaneous values of the narrowband gaussian noise are only very slightly correlated with one another, to be the distance,

$$y_{\mathrm{coh}} = \frac{1}{2f_{\mathrm{BW}}}, \tag{A6}$$

to the first zero in the sinc window. Note that the coherence length is the reciprocal of the full linear bandwidth.

The squared autocorrelation function, $|\rho(y)|^2$, quantifies the proportion of variance that a measurement at location y_0 accounts for at location $y_0 + y$. It is plotted in Figure 27. Beyond the coherence length, the squared correlation is very low; it never exceeds 0.05. Observations of instantaneous values of the noise that are taken at distances of more than a coherence length are thus approximately independent. A stimulus that extends over several coherence lengths thus offers several independent observations of the noise intensity, and an ideal observer will exploit this information.

Correlation Within a Coherence Length

The squared autocorrelation is not uniformly high even within a coherence length. There are zeros in the autocorrelation, whenever either the sinc term or the cosine term vanish. The first zero within a coherence length occurs when the cosine term first vanishes, at $y_1 = 1/(4f_a)$. Two instantaneous measurements of the noise taken at a distance of y_1 are thus independent, and they thus account for independent proportions of variance across the stimulus.

According to Eq. A6, a 1/3-octave gaussian noise centered at 0.2 cpd has a coherence length of 21.6°. The 13.4° stimuli used in all five experiments thus extended over only about 0.62 of this noise's coherence length. Two instantaneous measurements, taken at a distance of y_1 apart from one another, account for

$$\frac{1}{y_{coh}} \int_{-0.62 \cdot y_{coh}/2}^{0.62 \cdot y_{coh}/2} \left[|\rho(y)|^2 + |\rho(y+y_1)|^2 \right] dy \approx 0.90 \qquad \text{(A7)}$$

of the variance in this interval.

Appendix B

Model of Psychometric Functions for Detection

Herein, a near-ideal observer model is developed of the psychometric function for detection of a random-phase sinusoid in the presence of narrowband gaussian noise. It is assumed that the observation interval extends no farther than 0.62 of a coherence length (see Appendix A).

According to Eq. A7, two instantaneous measurements of a 1/3-octave gaussian noise, appropriately spaced, account for 90% of the variance in 0.62 of a coherence length. If the observation interval were narrower than 0.62 of a coherence length, these two measurements would account for an even greater proportion of the variance. A near-ideal model of detection of a random-phase sinusoid in the presence of narrowband noise can thus be developed by assuming that only these two independent measurements are conducted. This is not in any way meant to imply that human subjects actually perform the task in this way. The purpose of the model is to describe the psychometric function of a near-ideal observer, against which human performance can be judged.

Because the sinusoidal signal is of random phase, a phase-insensitive model must be developed. An energy detector is such a model. In this model, the decision variable is a function of the sum of the squared amplitudes of the measurements.

Let N denote the long-term noise intensity. Also, let A_1 and A_2 denote the two independent instantaneous amplitude measurements of the noise-alone stimulus. A_1 and A_2 are each gaussian random variables with mean zero and variance N. The quantity

$$\frac{A_1^2 + A_2^2}{N} \tag{B1}$$

is a χ^2 random variable with 2 degrees of freedom (df).

Let S denote the signal intensity (This implies that the signal *amplitude* is $\sqrt{2S}$).

Also, let B_1 and B_2 denote the two instantaneous amplitude measurements of the signal-plus-noise stimulus. B_1 and B_2 are each gaussian random variables with variance N. The effect of the signal is to add offsets to the measurements. Thus, the mean of B_1 is

$$\mu_{B_1} = \sqrt{2S}\cos(2\pi f_g y + \varphi), \tag{B2}$$

and the mean of B_2 is

$$\begin{aligned}\mu_{B_2} &= \sqrt{2S}\cos\left[2\pi f_g (y - y_1) + \varphi\right] \\ &= \sqrt{2S}\cos\left[2\pi f_g y - \frac{f_g}{f_a}\frac{\pi}{2} + \varphi\right],\end{aligned} \tag{B3}$$

where f_g, f_a, and y_1 are defined in Appendix A respectively as the logarithmic center frequency of the noise, the arithmetic center frequency of the noise, and the distance to the first zero in the autocorrelation function, and φ is the (random) phase angle of the sinusoidal signal. For narrowband noises, $f_g / f_a \approx 1$, and, therefore,

$$\mu_{B_2} \approx \sqrt{2S}\sin\left(2\pi f_g y + \varphi\right). \tag{B4}$$

The quantity

$$\frac{B_1^2 + B_2^2}{N} \tag{B5}$$

is thus a noncentral chi-squared random variable (denoted χ'^2) with 2 df and noncentrality parameter

$$\lambda = \frac{\mu_{B_1}^2 + \mu_{B_2}^2}{2N} \approx \frac{S}{N}. \tag{B6}$$

Note that the random phase angle, φ, of the signal does not appear in Eq. B6. This is because μ_{B_1} and μ_{B_2} are approximately in quadrature ($\pi/2$) phase; the sums of their squares is thus approximately 1, regardless of the shared offset in their phase angles. Even if the restriction that $f_g / f_a \approx 1$ (narrowband noises) is lifted, the expectation of λ across

all values of the random phase angle remains $E_{\varphi}[\lambda] = S/N$.

The observer picks as her guess for the signal interval that interval with the higher energy. Thus, if the ratio

$$F' = \frac{B_1^2 + B_2^2}{A_1^2 + A_2^2} = \frac{(B_1^2 + B_2^2)/N}{(A_1^2 + A_2^2)/N} \tag{B7}$$

exceeds unity, then the observer chooses the correct interval as the signal interval, and otherwise chooses the incorrect interval. The probability of a correct response is thus $P(F' > 1)$. F' is the ratio of a χ'^2 random variable with 2 df and noncentrality parameter λ, to a χ^2 random variable with 2 df. F' is thus a *noncentral F-ratio* random variable with 2 numerator df, 2 denominator df, and noncentrality parameter λ. Although the cumulative distribution function, $F'(x;2,2,\lambda)$, for such a random variable is generally complicated, it has the simple closed form

$$F'(1;2,2,\lambda) = \frac{1}{2}e^{-\lambda/2} \tag{B8}$$

for the special case $x = 1$. Thus,

$$\begin{aligned} P(F' > 1) &= 1 - F'(1;2,2,S/N) \\ &= 1 - \frac{1}{2}\exp\left(\frac{-S}{2N}\right) \end{aligned} \quad , \tag{B9}$$

which is algebraically equivalent to the Quick (1974) psychometric function

$$\begin{aligned} P(\text{correct response}) &= 1 - \frac{1}{2}2^{-S/(2N\ln 2)} \\ &\approx 1 - \frac{1}{2}2^{-0.72S/N} \end{aligned} \tag{B10}$$

Appendix C

Model of Psychometric Functions for SF Discrimination

Herein, a model is developed of the psychometric function for the discrimination between two sinusoids of different SF, equal intensities, and independent random phases.

Assume that two sinusoids of intensity I, one of frequency f and the other of slightly higher frequency $f\delta$ $(\delta > 1)$, are passed by a trapezoidal filter's high-frequency skirt (the analysis for the low-frequency skirt is nearly identical). If filter has center frequency ν , passband parameter p, and skirt parameter s, then, according to Eq. 6, the output of the filter in response to the lower-SF sinusoid is

$$\left(\frac{f}{2^p \nu}\right)^{-1/s}, \tag{C1}$$

and the output of the filter in response to the higher-SF sinusoid is

$$\left(\frac{f\delta}{2^p \nu}\right)^{-1/s}. \tag{C2}$$

The ratio of the outputs is thus $\delta^{1/s}$. The proportion correct can therefore be expected to vary monotonically with $\delta^{1/s}$.

Assume that the particular form of the psychometric function is

$$\begin{aligned} P(\text{correct response}) &= 1-\frac{1}{2}\exp\left(-a\ln\delta^{\left(s^{-\eta}\right)}\right), \\ &= 1-\frac{1}{2}\delta^{-s^{-\eta}a} \end{aligned} \tag{C3}$$

where a and η are free parameters greater than zero. Eq. C3 represents a modification of the Quick (1974) psychometric function of the sort derived in Appendix B.

Setting P(correct response)=0.75, we find that the corresponding (threshold) value

of δ is

$$\delta_{\text{thresh}} = 2^{\left(s^{\eta}/a\right)}. \tag{C4}$$

Because

$$\ln\left(\frac{\Delta f}{f}\right) = \ln\left(\frac{f\delta - f}{f}\right), \tag{C5}$$
$$= \ln(\delta - 1)$$

the value of $\ln(\Delta f/f)$ at threshold is

$$\ln\left(\frac{\Delta f}{f}\Big|_{\text{thresh}}\right) = \ln\left[-1 + 2^{\left(s^{\eta}/a\right)}\right], \tag{C6}$$

which is Eq. 13.

The particular psychometric function displayed in Eq. C3 was posited because, for the ranges of parameters encountered in the present paper model, the steepness of this psychometric function is approximately invariant. More precisely, psychometric functions of the form in Eq. C3, whose parameters have different values, are approximately simply shifted on the $\ln(\Delta f/f)$ axis. The advantage of positing such a functional form is that the particular definition of threshold--in our case 75% correct--is unimportant. This is particularly attractive given that the threshold measurements in Experiment 5 are based on the mean of staircase turnarounds, and the proportion correct to which the present experiments' staircase procedures converge is difficult to derive. However, is important to note that the threshold equation, Eq. C4, is not tightly bound to this particular psychometric function. The same threshold equation can be derived from other reasonable psychometric functions. Because the raw data in Experiment 5 were not fit with the specific psychometric function in Eq. C4, the threshold estimates can be expected to be robust with respect to this assumption.

Table 1. Proportions of variance accounted for by unity-slope lines in dB-dB plots of threshold versus masker intensity.

	Subject		
Signal SF (cpd)	AC	AG	DR
0.2	.993	.972	.970
0.2828	.993	.808	
0.4	.972	.923	
0.5657	.863	.840	
0.8	.982	.963	.895
1.1314	.982	.951	
1.6	.961	.931	

Table 2. Best-fitting parameter values for trapezoidal filter to subject AC's notched and bandpass noise data at each signal frequency. Also listed are pooled root mean squared error (rmse), rmse for notched data alone, and rmse for bandpass data alone, as well as estimated 3 dB octave bandwidths.

sig freq (cpd)	0.2	0.28	0.4	0.57	0.8	1.13	1.6
s (oct)	0.25	0.29	0.29	0.26	0.31	0.63	0.52
p (oct)	0.42	0.32	0.33	0.34	0.70	0.63	0.18
rmse (dB)	2.64	3.02	3.20	4.67	2.83	1.99	2.03
$rmse_{notch}$ (dB)	1.72	2.73	2.88	4.80	3.07	2.30	1.93
$rmse_{bp}$ (dB)	4.62	3.70	4.03	5.49	2.93	2.48	2.18
$2s$ (oct)	0.50	0.59	0.58	0.51	0.62	1.25	1.03
$2p$ (oct)	0.84	0.61	0.67	0.69	1.39	1.26	0.35
$BW_{3\ dB}$ (oct)	1.34	1.23	1.25	1.21	2.01	1.89	1.37

Table 3. Best-fitting parameter values for trapezoidal filter to subject AG's notched and bandpass noise data at each signal frequency. Also listed are pooled root mean squared error (rmse), rmse for notched data alone, and rmse for bandpass data alone, as well as estimated 3 dB octave bandwidths.

sig freq (cpd)	0.2	0.28	0.4	0.57	0.8	1.13	1.6
s (oct)	0.17	0.25	0.26	0.24	0.37	0.26	0.61
p (oct)	0.61	0.14	0.11	0.25	0.40	0.78	0
rmse (dB)	5.87	4.43	3.90	4.10	4.89	7.87	3.82
$rmse_{notch}$ (dB)	4.14	5.14	3.69	5.09	5.75	6.15	3.02
$rmse_{bp}$ (dB)	7.19	3.56	5.37	3.35	3.84	9.58	4.49
2*s* (oct)	0.34	0.49	0.51	0.49	0.74	0.53	1.21
2*p* (oct)	1.22	0.27	0.22	0.50	0.80	1.55	0
$BW_{3\ dB}$ (oct)	1.56	0.76	0.73	0.99	1.54	2.08	1.21

Table 4. Best-fitting parameter values for trapezoidal filter to subject DR's notched and bandpass noise data at 0.2 cpd, the SF where he provided a complete data set. Also listed are pooled root mean squared error (rmse), rmse for notched data alone, rmse for bandpass data alone, as well as estimated 3 dB octave bandwidths.

sig freq (cpd)	0.2
s	0.28
p	0.36
rmse (dB)	4.26
$rmse_{notch}$ (dB)	3.93
$rmse_{bp}$ (dB)	4.57
$2s$ (oct)	0.55
$2p$ (oct)	0.72
$BW_{3\ dB}$ (oct)	1.28

Table 5. Root mean squared errors (rmse) for trapezoidal filter model's predictions of subject AC's pooled, notched, and bandpass noise data at each signal frequency. Entries headed *free* refer to independent fits at each spatial frequency (14 df total); entries headed *constrained* refer to model predictions when parameters were fit as second-order polynomials in logarithmic SF (6 df total).

sig freq (cpd)	0.2	0.28	0.4	0.57	0.8	1.13	1.6	pooled across sig freq
rmse (dB)								
free	2.64	3.02	3.20	4.67	2.83	1.99	2.03	3.03
constrained	3.57	3.31	3.26	5.32	3.17	4.63	3.10	3.85
$rmse_{notch}$ (dB)								
free	1.72	2.73	2.88	4.80	3.07	2.30	1.93	2.93
constrained	2.03	2.86	2.24	5.14	3.83	6.06	3.81	3.90
$rmse_{bp}$ (dB)								
free	4.62	3.70	4.03	5.49	2.93	2.48	2.18	3.12
constrained	4.62	3.70	4.03	5.49	2.93	2.47	2.18	3.80

Table 6. Root mean squared errors (rmse) for trapezoidal filter model's predictions of subject AG's pooled, notched, and bandpass noise data at each signal frequency. Entries headed *free* refer to independent fits at each spatial frequency (14 df total); entries headed *constrained* refer to model predictions when parameters were fit as second-order polynomials in logarithmic SF (6 df total).

sig freq (cpd)	0.2	0.28	0.4	0.57	0.8	1.13	1.6	pooled across sig freq
rmse (dB)								
free	5.87	4.43	3.90	4.10	4.89	7.87	3.82	5.16
constrained	8.23	6.23	6.24	5.62	4.97	8.11	5.13	6.48
$rmse_{notch}$ (dB)								
free	4.14	5.14	3.69	5.09	5.75	6.15	3.02	4.83
constrained	4.08	6.89	6.69	6.34	5.83	5.33	3.01	5.61
$rmse_{bp}$ (dB)								
free	7.19	3.56	5.37	3.35	3.84	9.58	4.49	5.82
constrained	10.91	5.48	5.76	4.80	3.93	10.15	6.60	7.25

Table 7. Best-fitting parameters and resulting root mean squared error (rmse) in fitting trapezoidal filter model to Experiment 4's absolute threshold data.

	Subject AC	Subject DR
N_1	0.273	0.257
ε	2.73	2.41
rmse (dB)	0.70	1.05

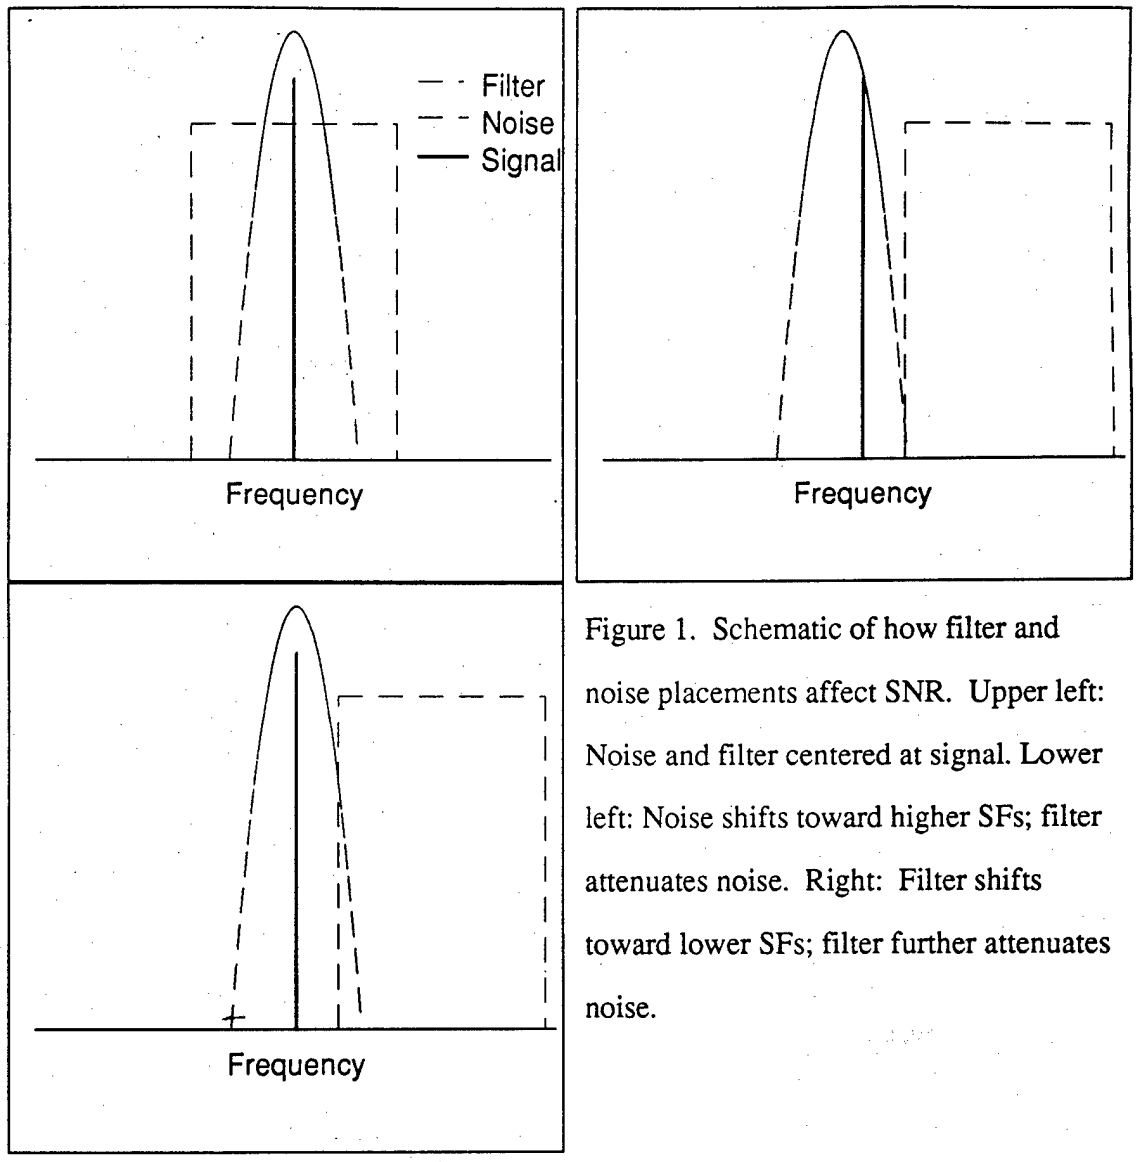

Figure 1. Schematic of how filter and noise placements affect SNR. Upper left: Noise and filter centered at signal. Lower left: Noise shifts toward higher SFs; filter attenuates noise. Right: Filter shifts toward lower SFs; filter further attenuates noise.

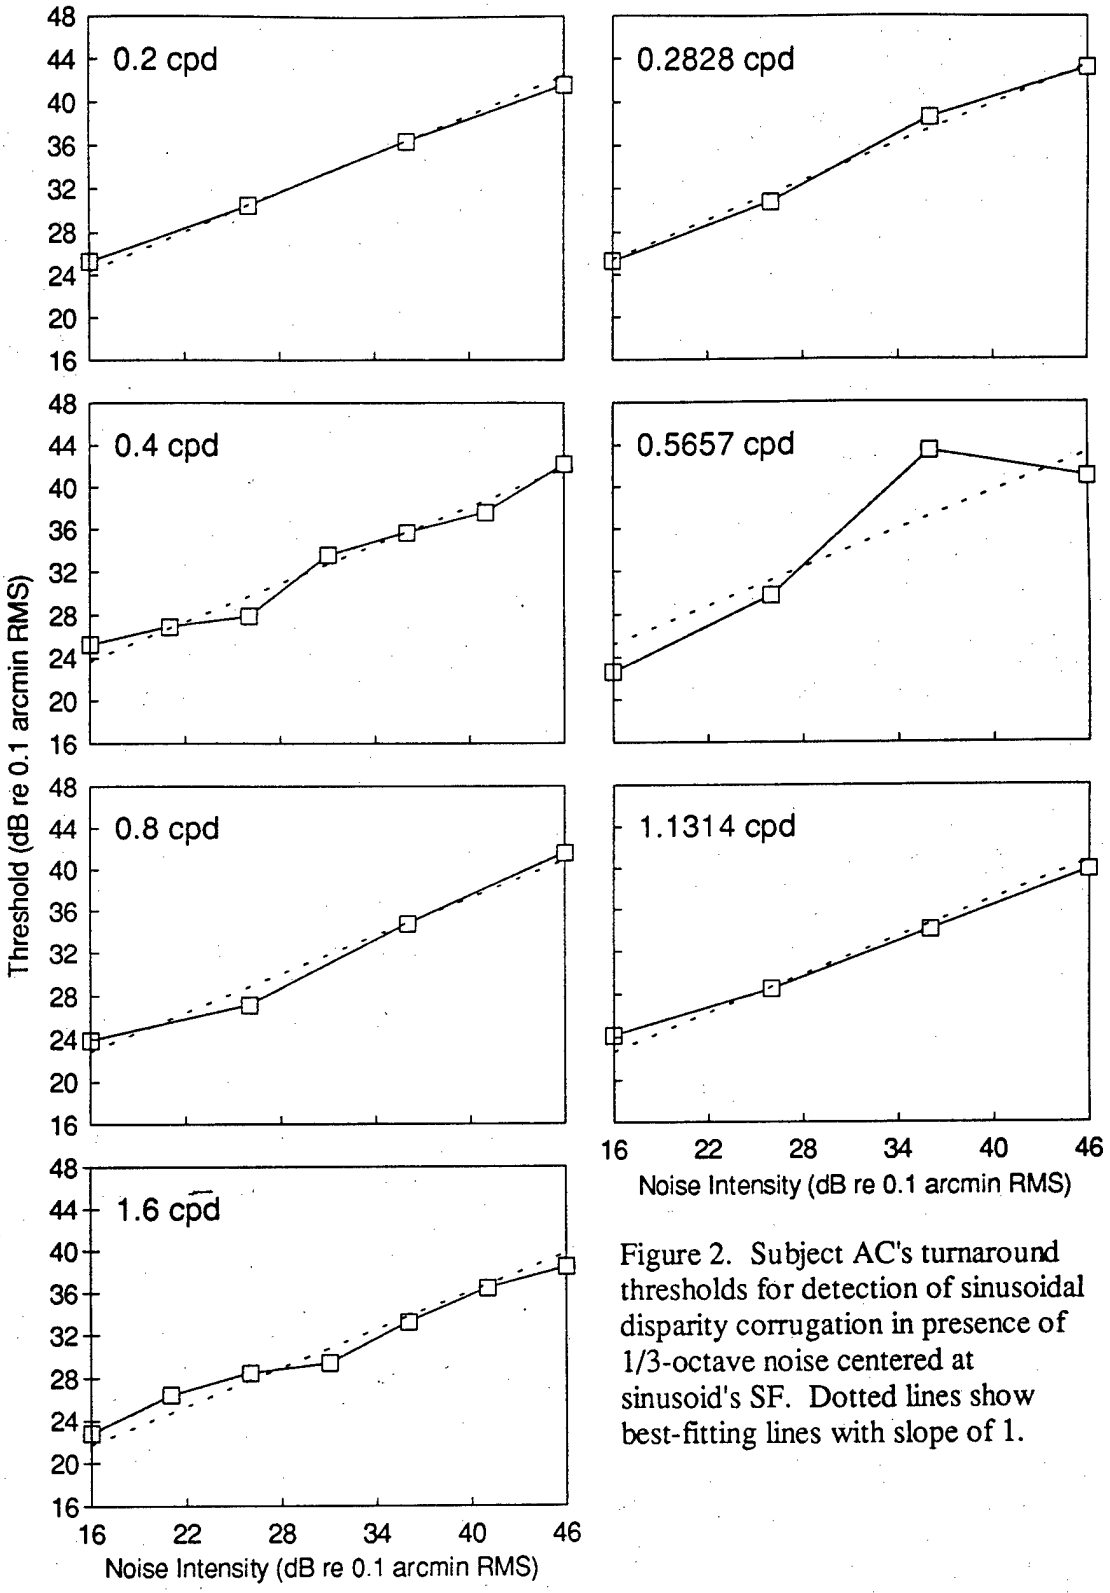

Figure 2. Subject AC's turnaround thresholds for detection of sinusoidal disparity corrugation in presence of 1/3-octave noise centered at sinusoid's SF. Dotted lines show best-fitting lines with slope of 1.

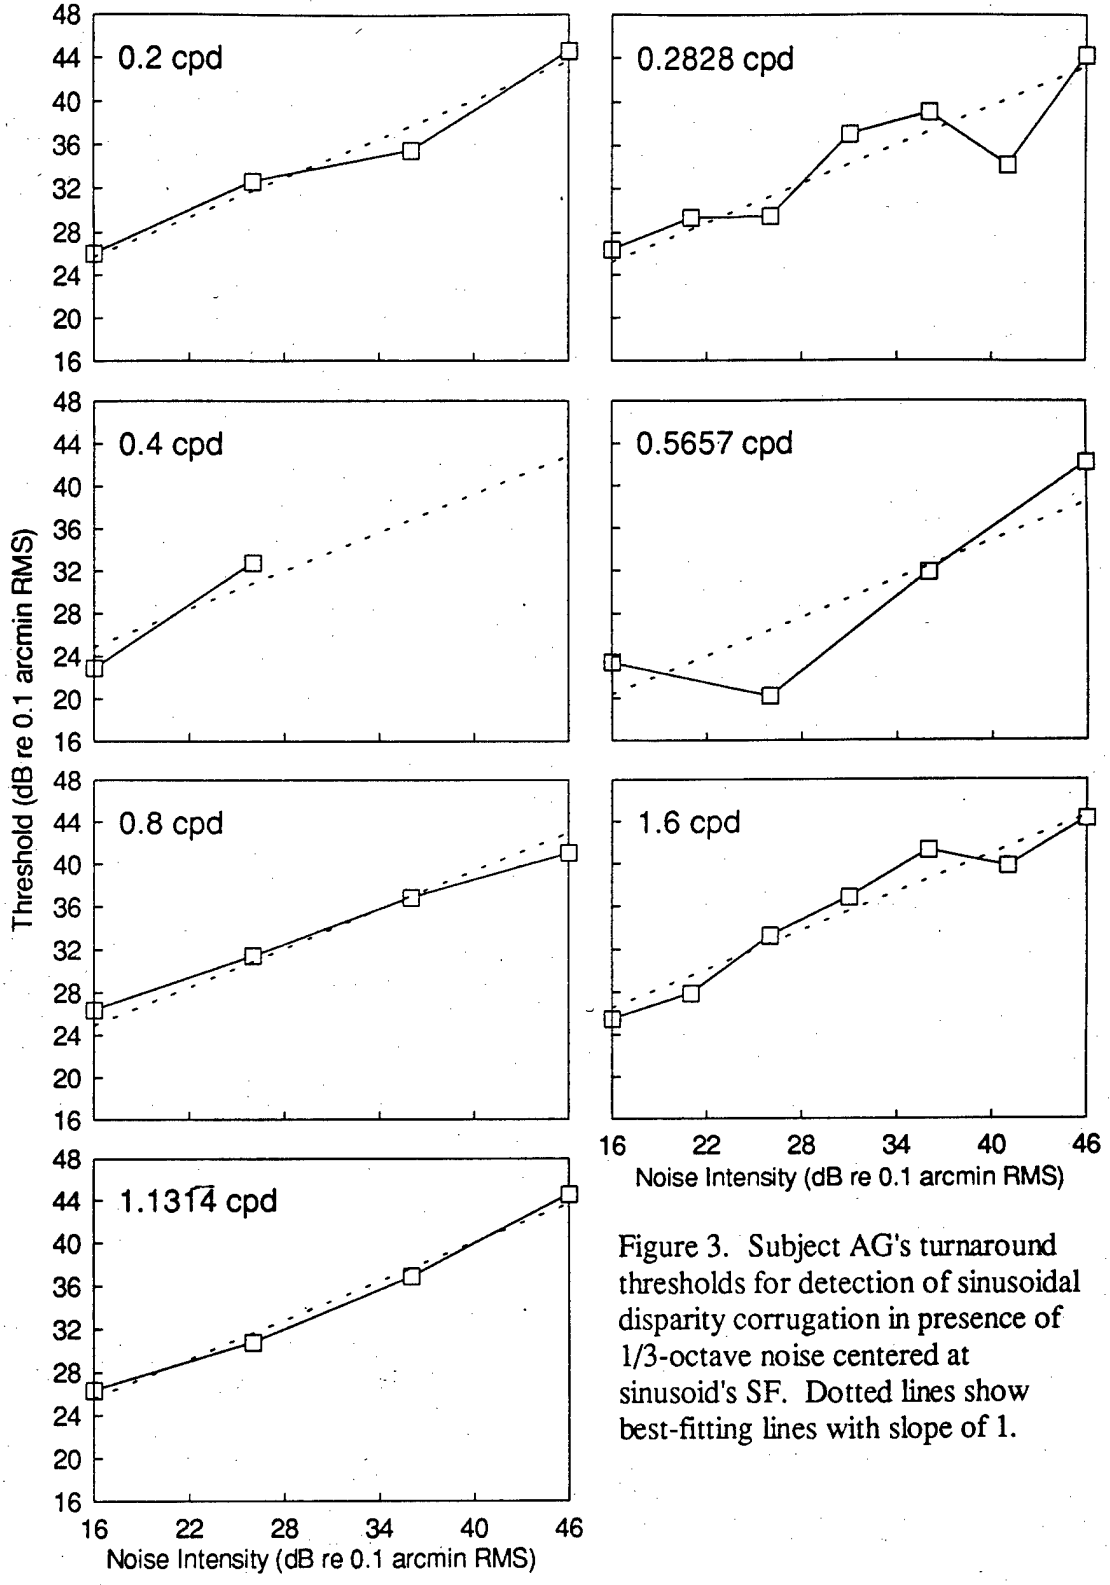

Figure 3. Subject AG's turnaround thresholds for detection of sinusoidal disparity corrugation in presence of 1/3-octave noise centered at sinusoid's SF. Dotted lines show best-fitting lines with slope of 1.

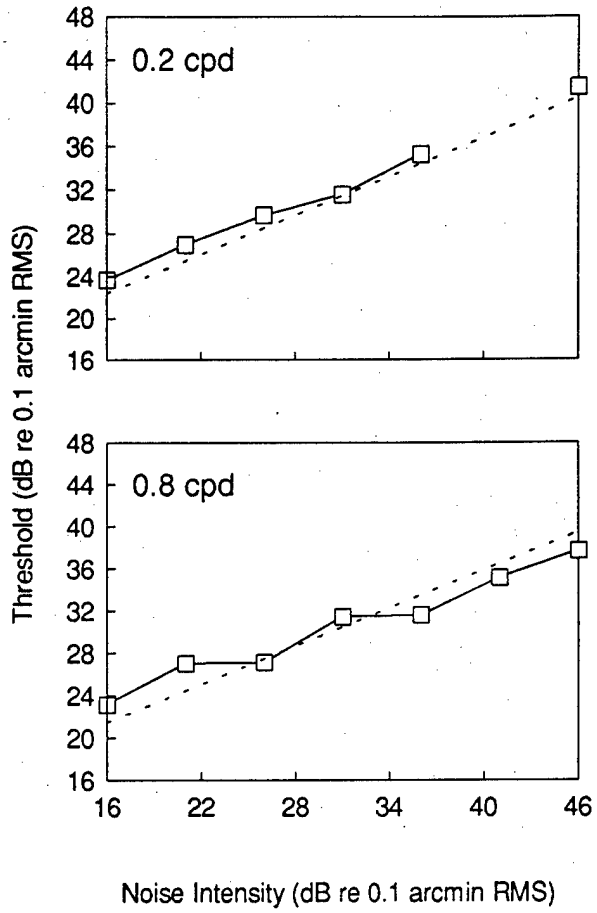

Figure 4. Subject DR's turnaround thresholds for detection of sinusoidal disparity corrugation in presence of 1/3-octave noise centered at sinusoid's SF. Dotted lines show best-fitting lines with slope of 1.

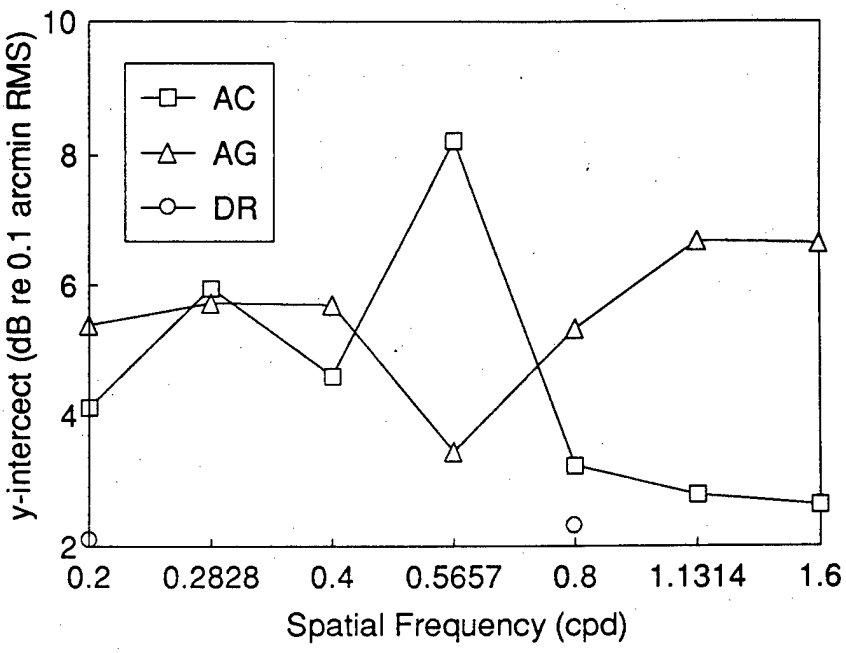

Figure 5. y-intercepts, as a function of spatial frequency, of unity-slope lines fit to threshold-versus-intensity plots for detection of sinusoidal disparity modulation in presence of 1/3-oct bandpass noise centered at sinusoid's spatial frequency.

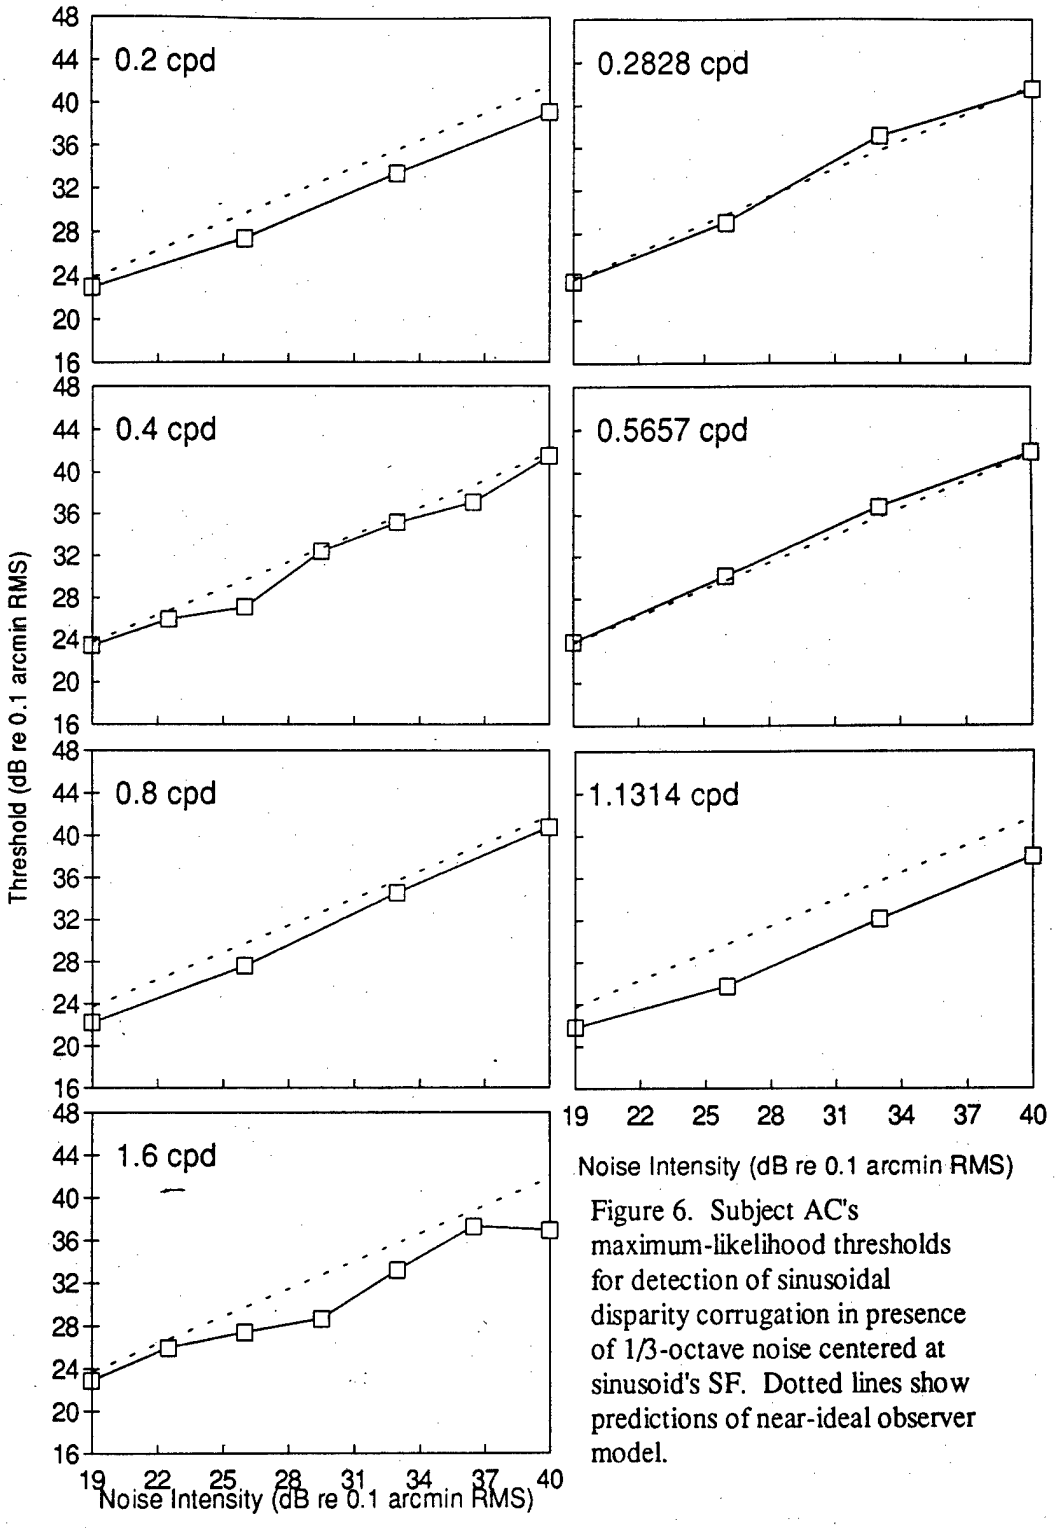

Figure 6. Subject AC's maximum-likelihood thresholds for detection of sinusoidal disparity corrugation in presence of 1/3-octave noise centered at sinusoid's SF. Dotted lines show predictions of near-ideal observer model.

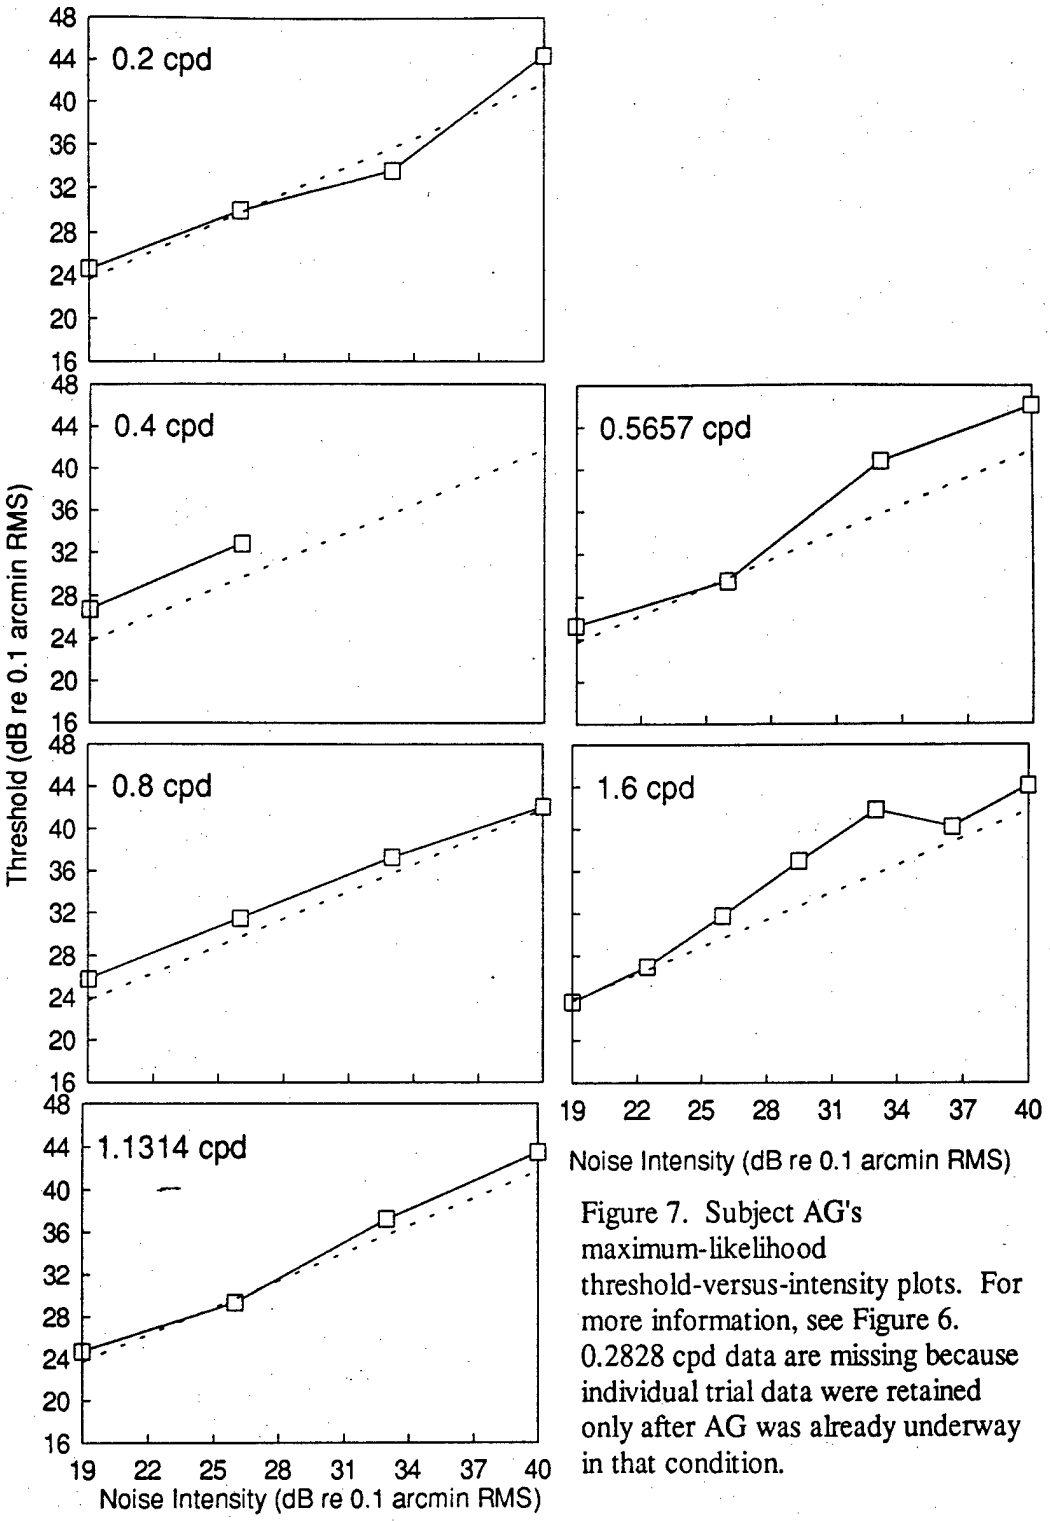

Figure 7. Subject AG's maximum-likelihood threshold-versus-intensity plots. For more information, see Figure 6. 0.2828 cpd data are missing because individual trial data were retained only after AG was already underway in that condition.

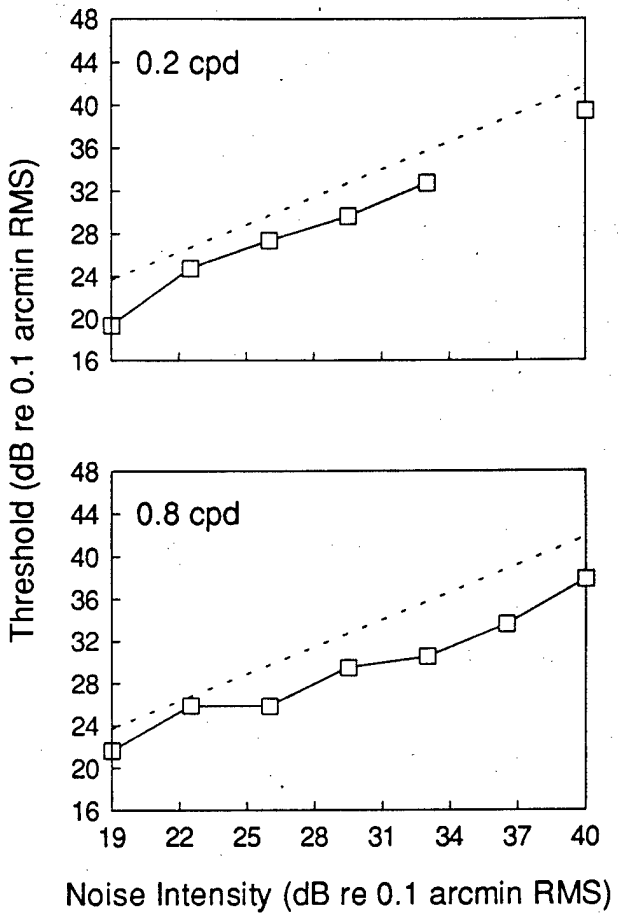

Figure 8. Subject DR's maximum-likelihood thresholds for detection of sinusoidal disparity corrugation in presence of 1/3-octave noise centered at sinusoid's SF. Dotted lines show predictions of near-ideal observer model.

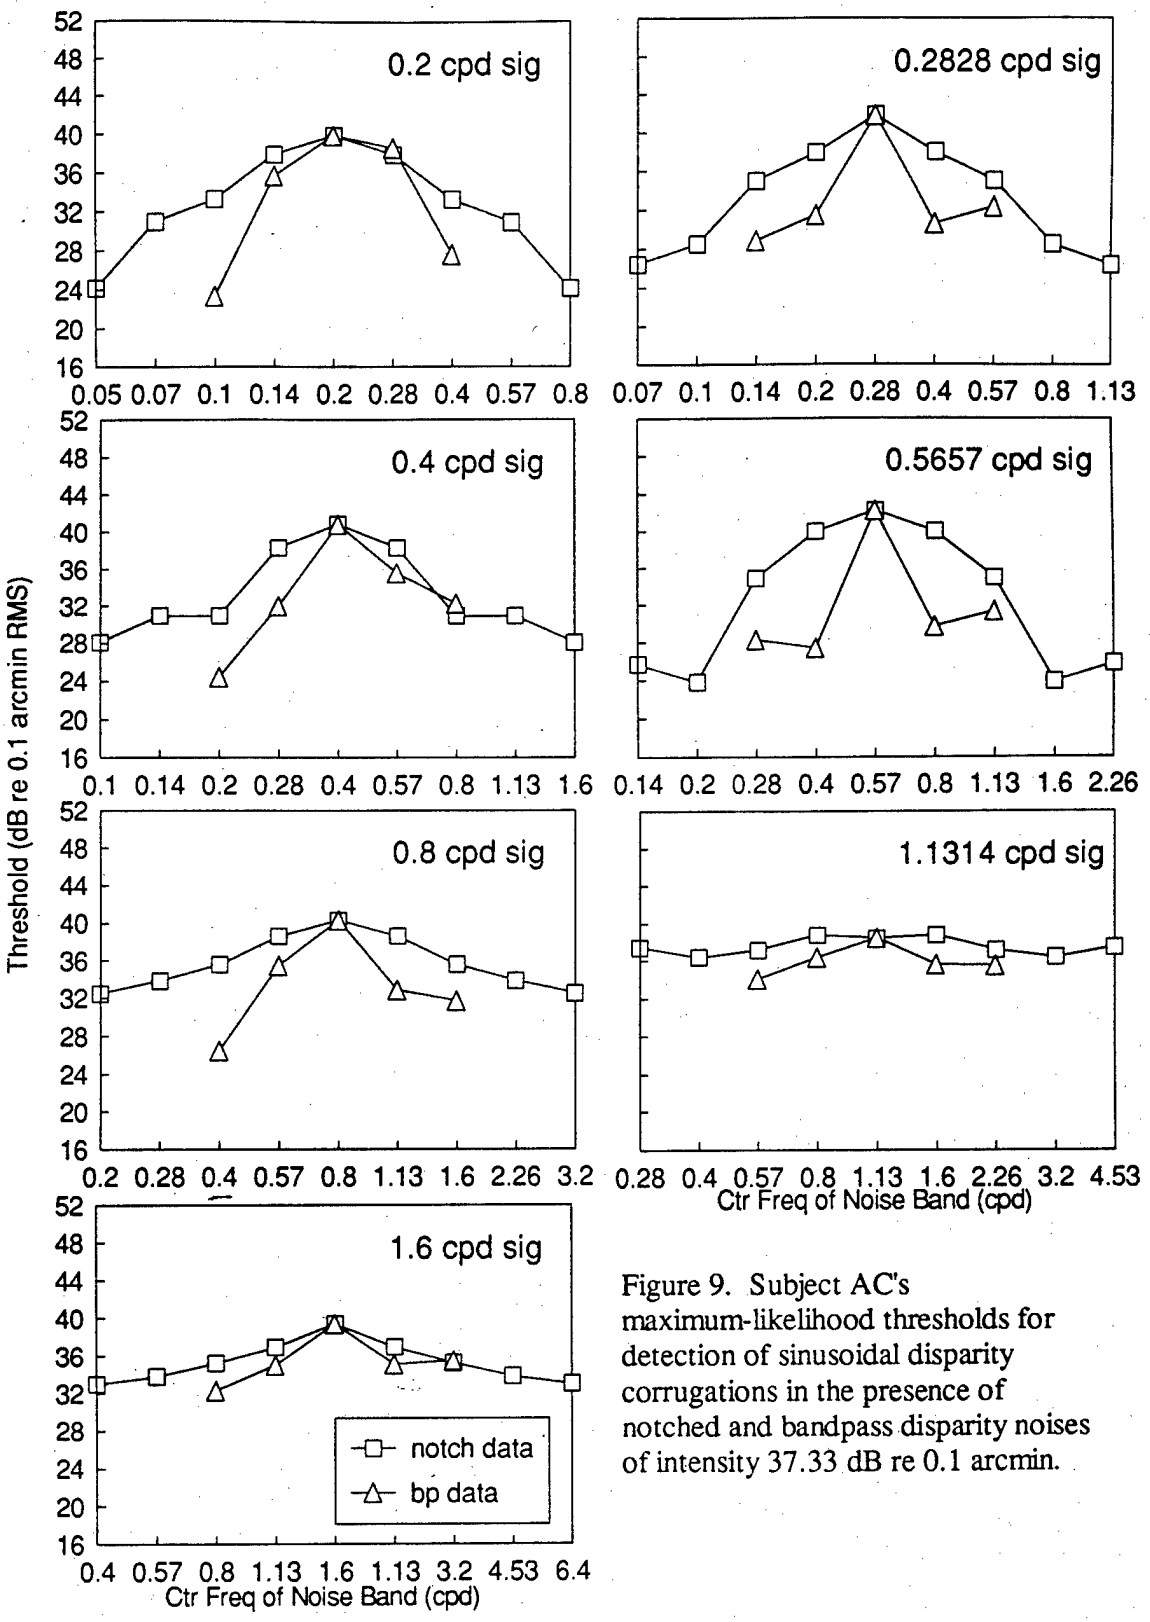

Figure 9. Subject AC's maximum-likelihood thresholds for detection of sinusoidal disparity corrugations in the presence of notched and bandpass disparity noises of intensity 37.33 dB re 0.1 arcmin.

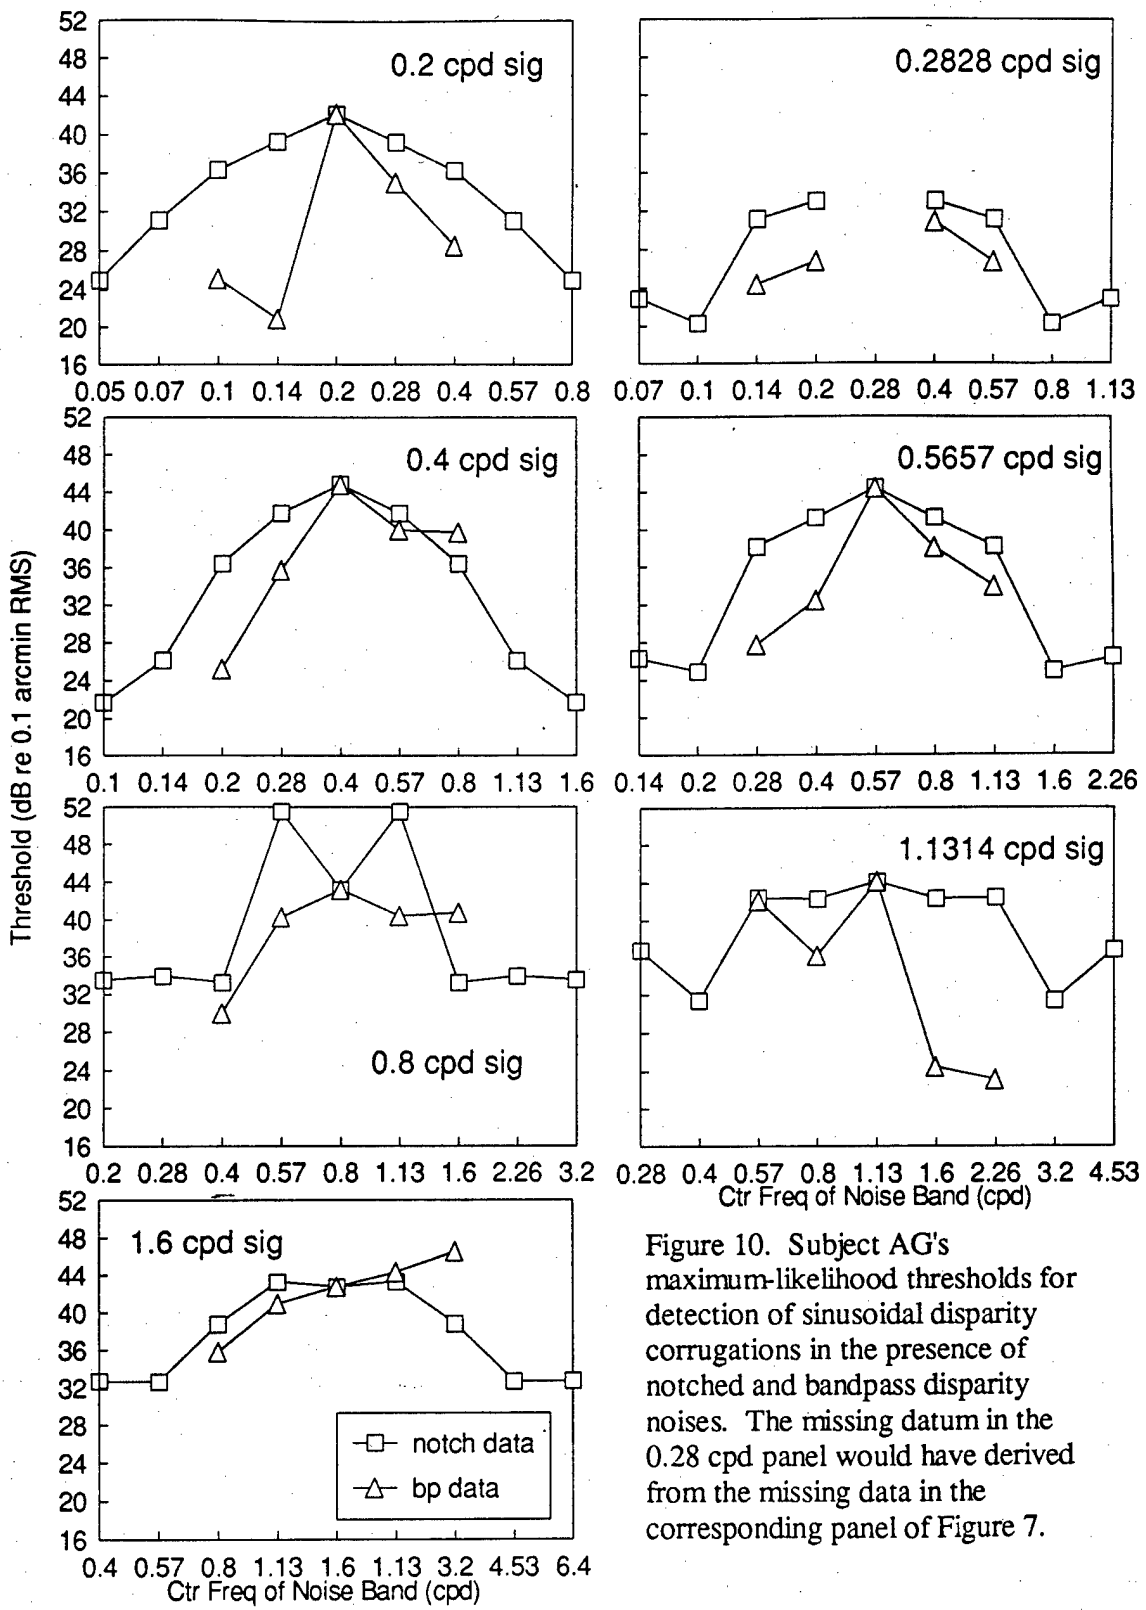

Figure 10. Subject AG's maximum-likelihood thresholds for detection of sinusoidal disparity corrugations in the presence of notched and bandpass disparity noises. The missing datum in the 0.28 cpd panel would have derived from the missing data in the corresponding panel of Figure 7.

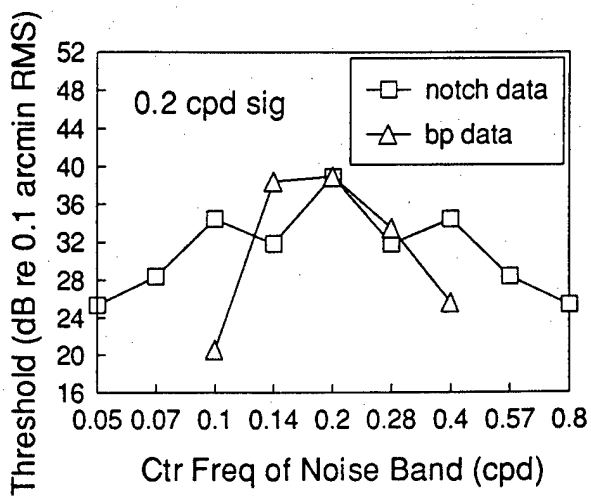

Figure 11. Subject DR's maximum-likelihood thresholds for detection of sinusoidal disparity corrugations in the presence of notched and bandpass disparity noises of intensity 37.33 dB re 0.1 arcmin.

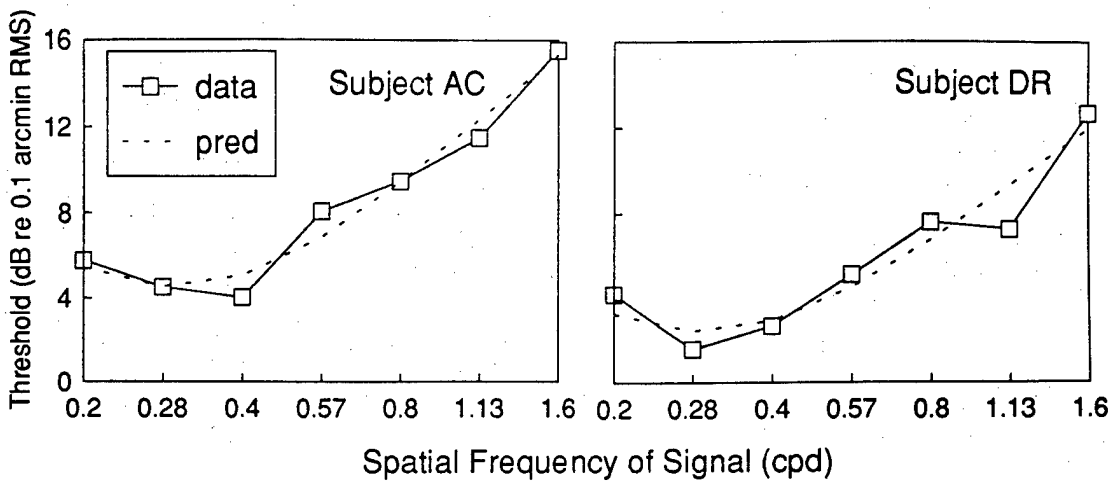

Figure 12. Maximum-likelihood thresholds for detection of sinusoidal disparity modulation in absence of noise.

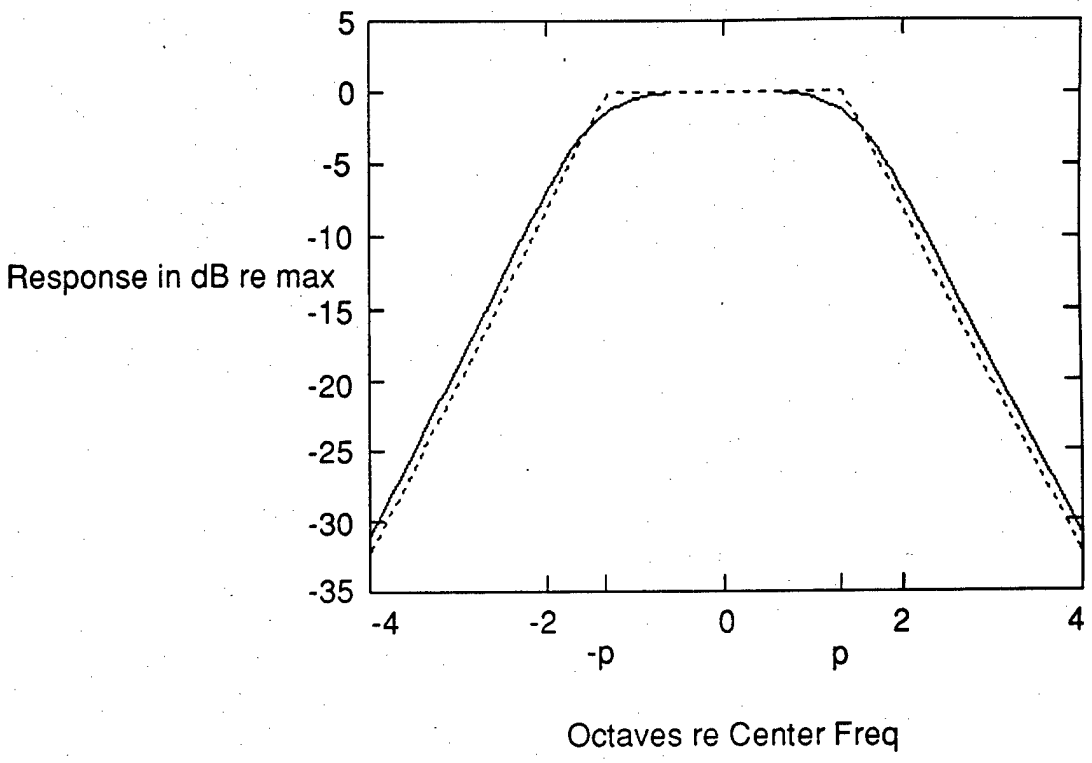

Figure 13. Butterworth bandpass filter (solid line) and "trapezoidal" filter (broken line). Butterworth filter is second order, with 3 dB bandwidth of 2 octaves. Trapezoid filter has parameters $s = 0.25$, $p = 1.32$.

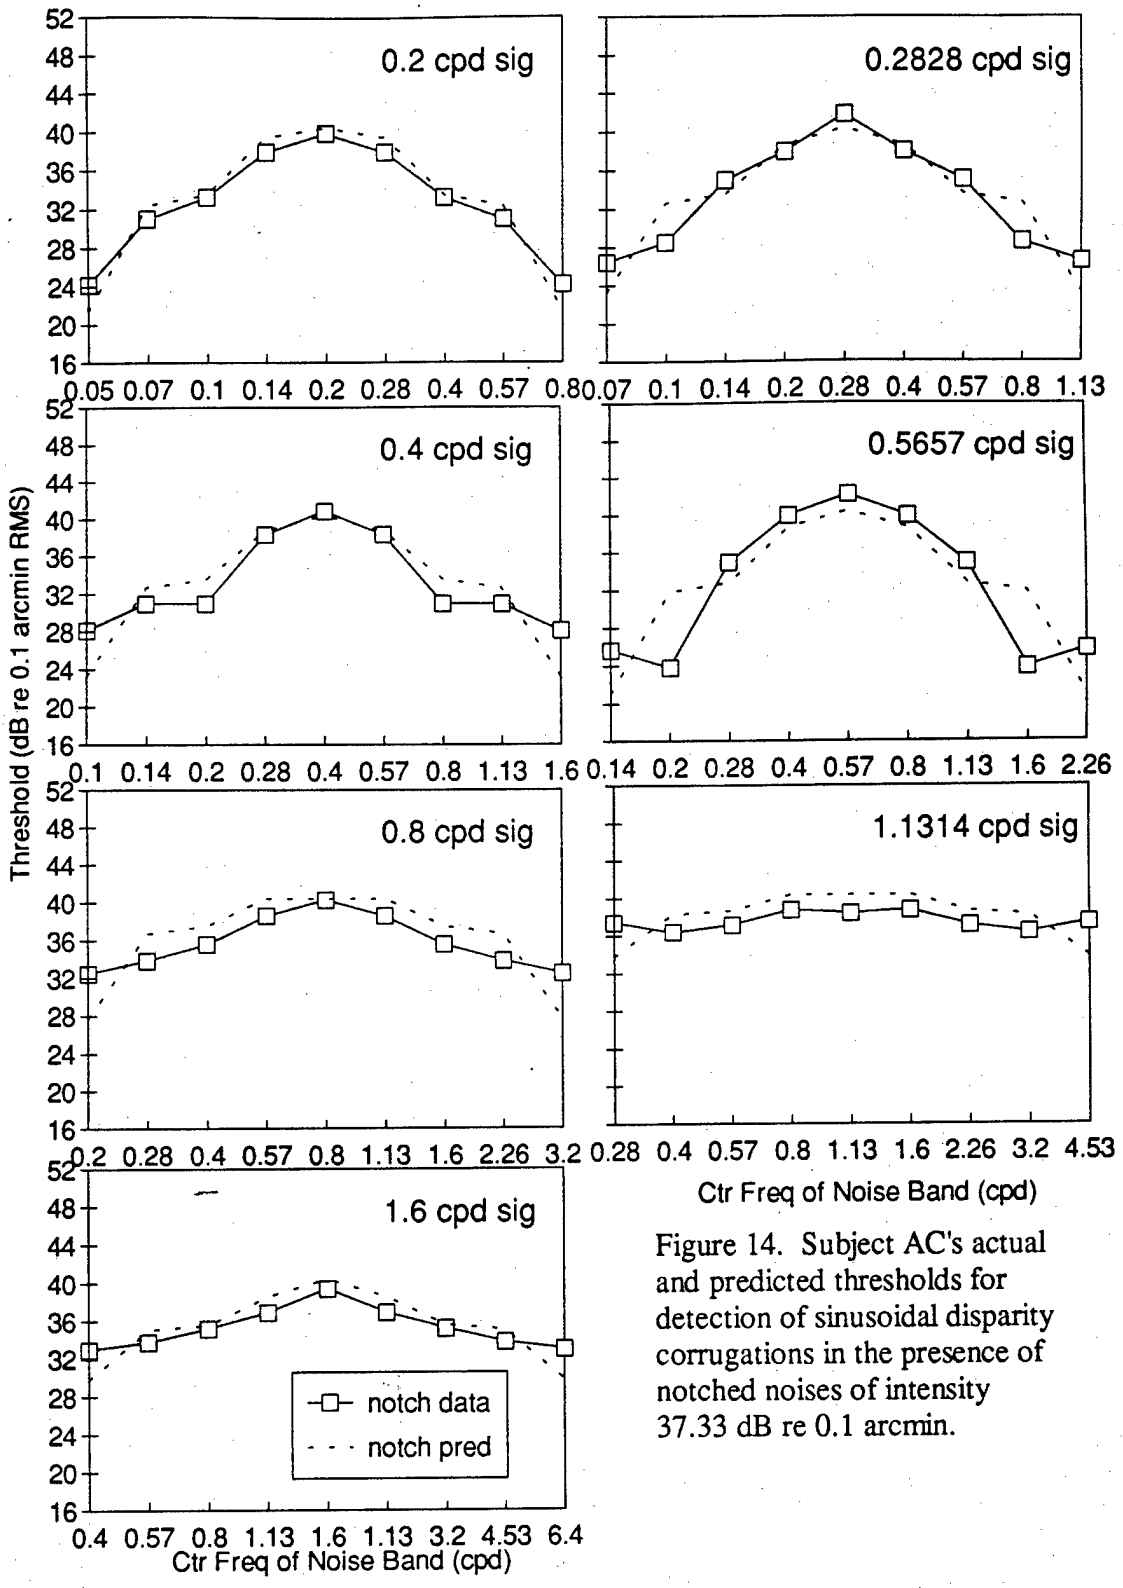

Figure 14. Subject AC's actual and predicted thresholds for detection of sinusoidal disparity corrugations in the presence of notched noises of intensity 37.33 dB re 0.1 arcmin.

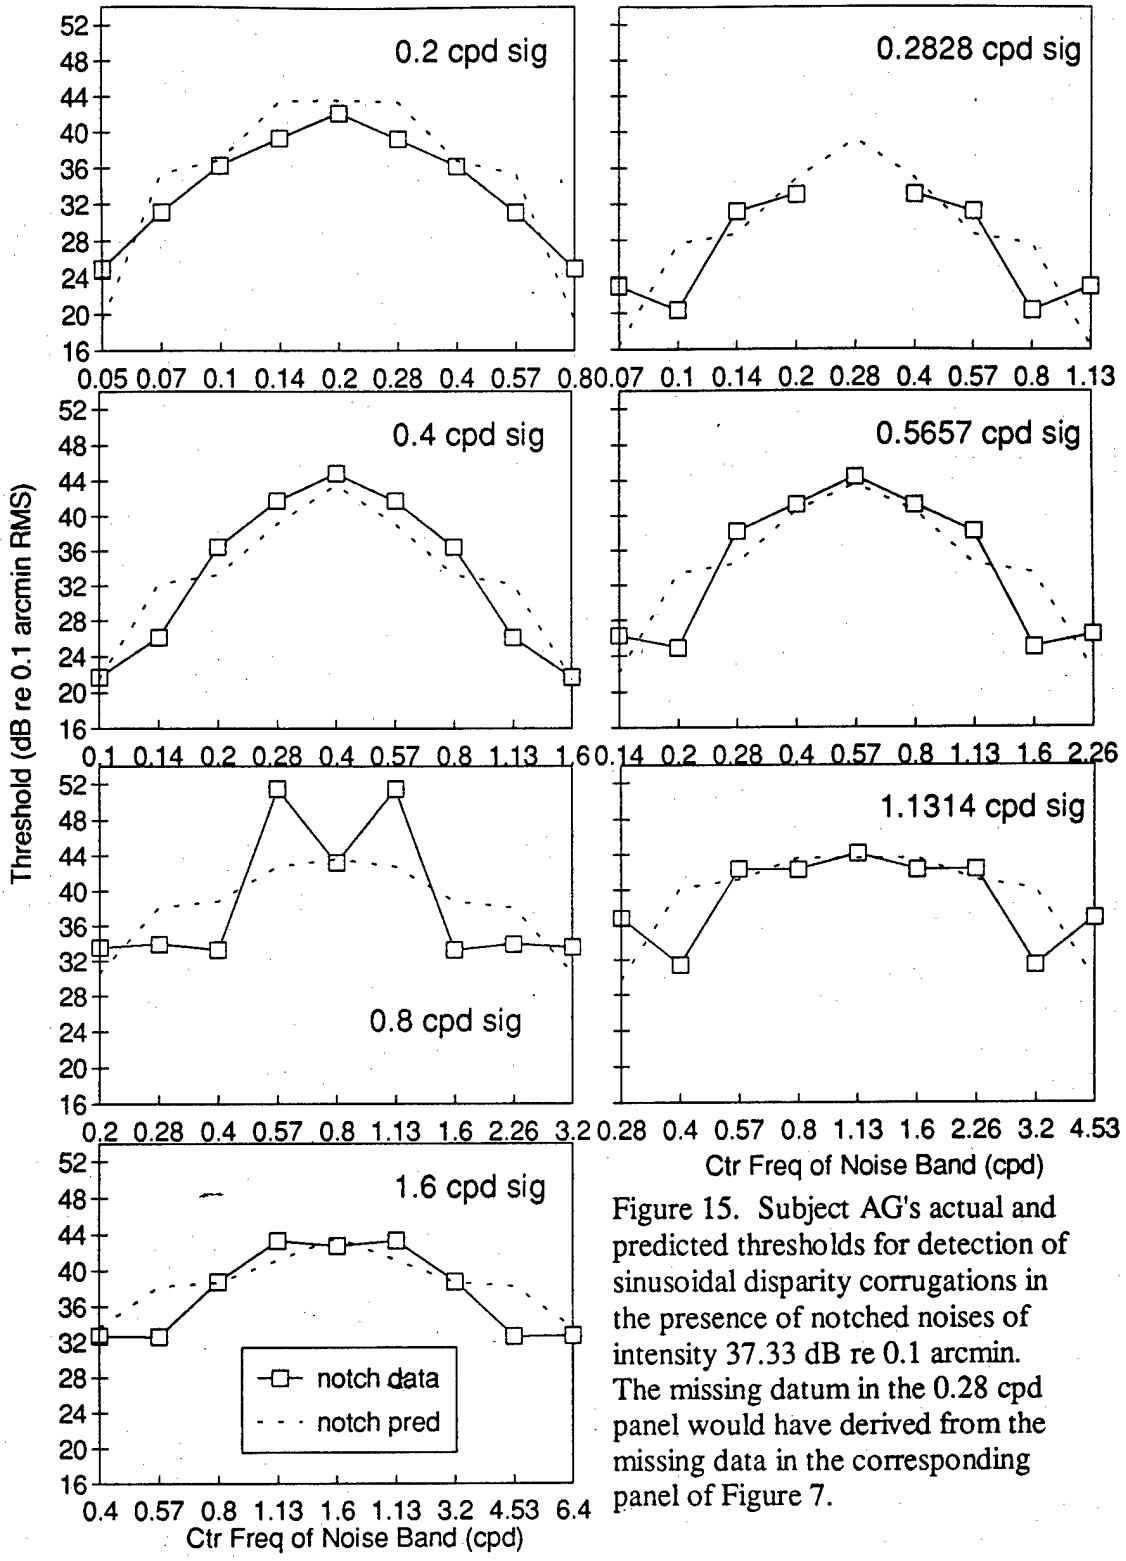

Figure 15. Subject AG's actual and predicted thresholds for detection of sinusoidal disparity corrugations in the presence of notched noises of intensity 37.33 dB re 0.1 arcmin. The missing datum in the 0.28 cpd panel would have derived from the missing data in the corresponding panel of Figure 7.

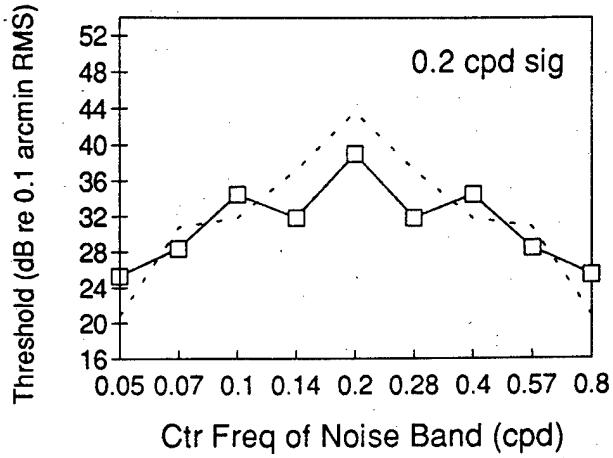

Figure 16. Subject DR's actual and predicted thresholds for detection of sinusoidal disparity corrugations in the presence of notched noises of intensity 37.33 dB re 0.1 arcmin.

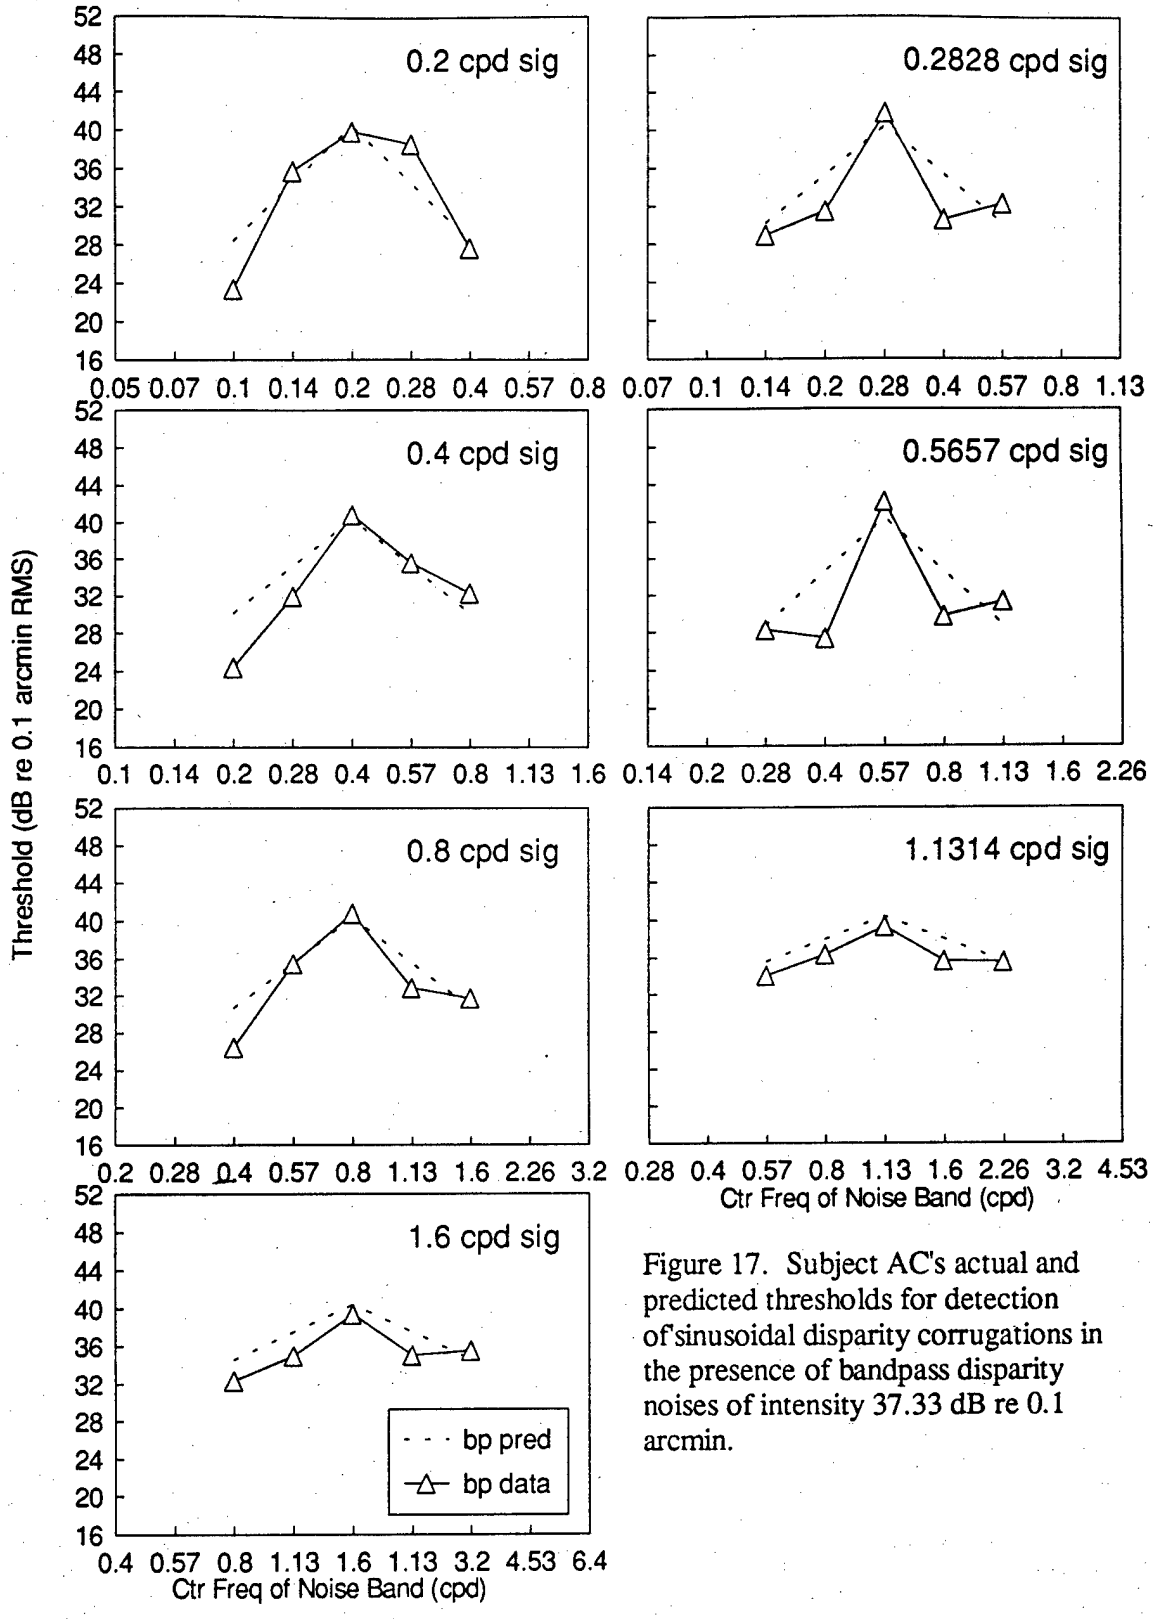

Figure 17. Subject AC's actual and predicted thresholds for detection of sinusoidal disparity corrugations in the presence of bandpass disparity noises of intensity 37.33 dB re 0.1 arcmin.

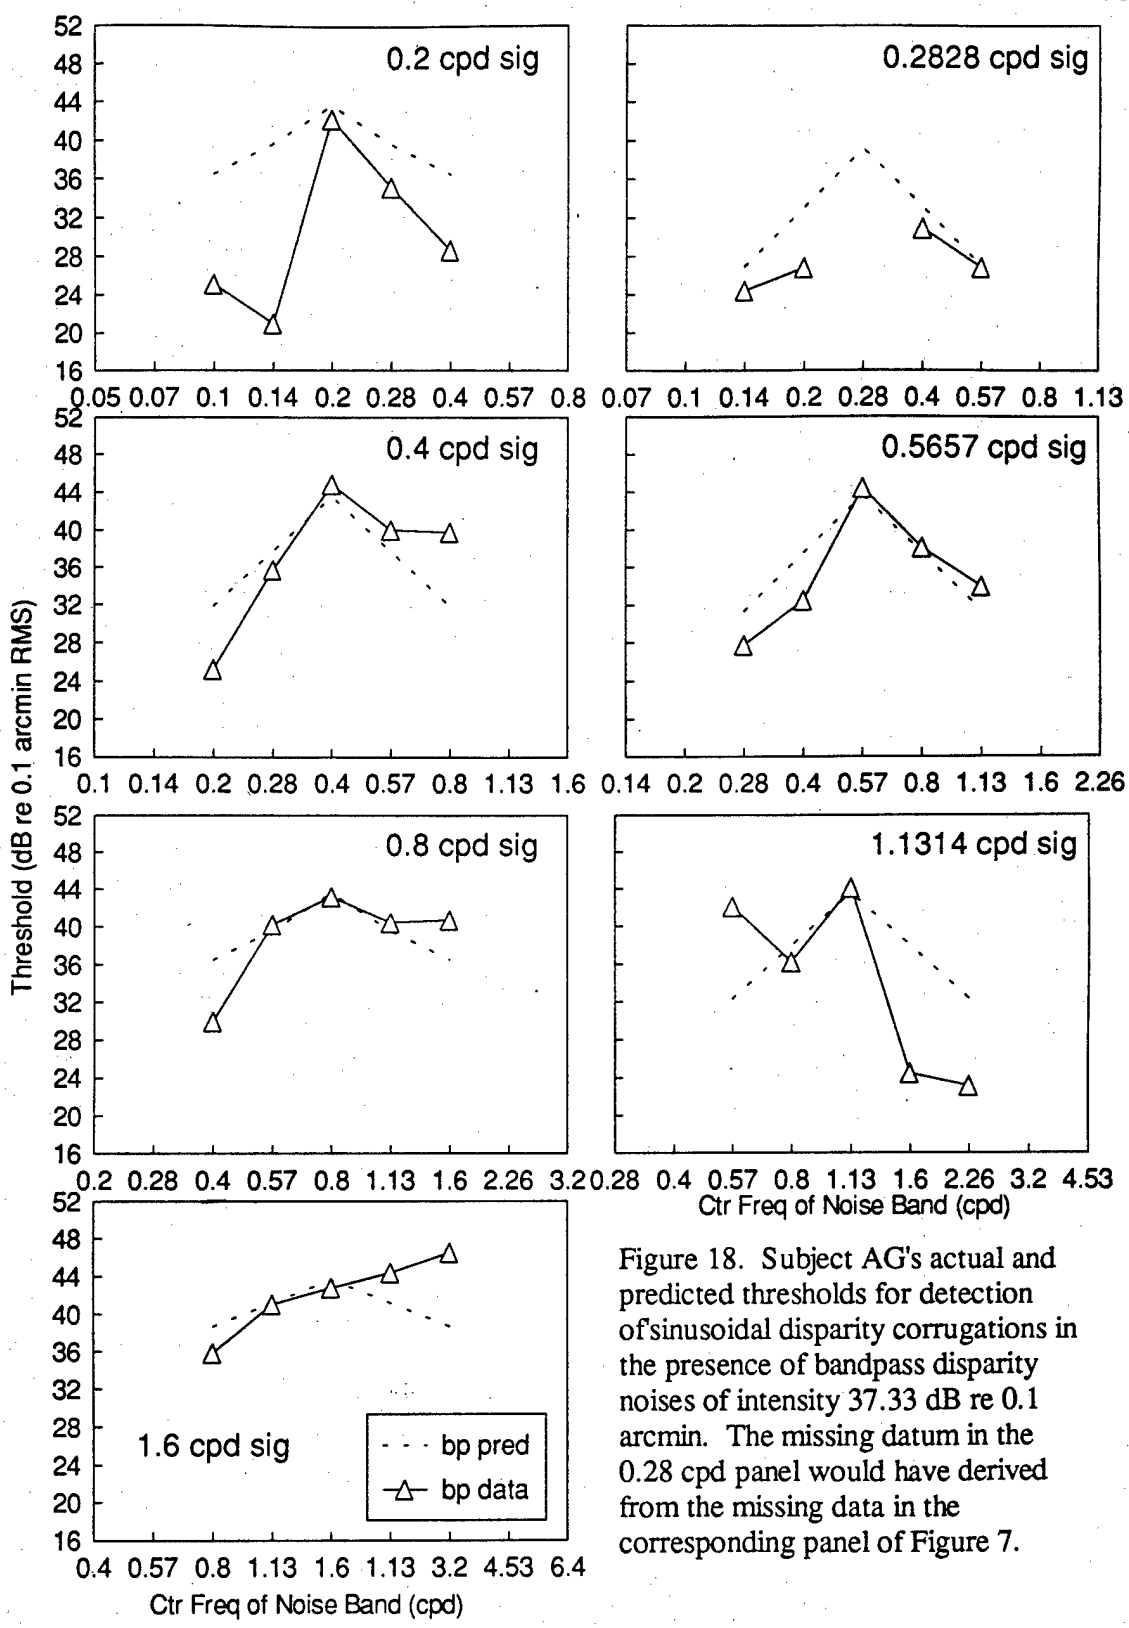

Figure 18. Subject AG's actual and predicted thresholds for detection of sinusoidal disparity corrugations in the presence of bandpass disparity noises of intensity 37.33 dB re 0.1 arcmin. The missing datum in the 0.28 cpd panel would have derived from the missing data in the corresponding panel of Figure 7.

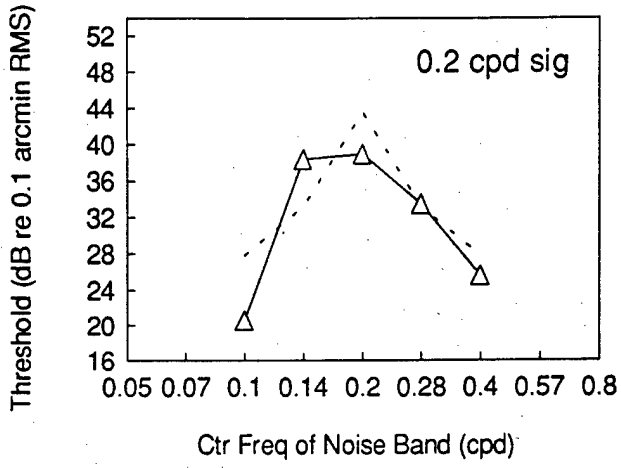

Figure 19. Subject DR's actual and predicted thresholds for detection of sinusoidal disparity corrugations in the presence of bandpass disparity noises of intensity 37.33 dB re 0.1 arcmin.

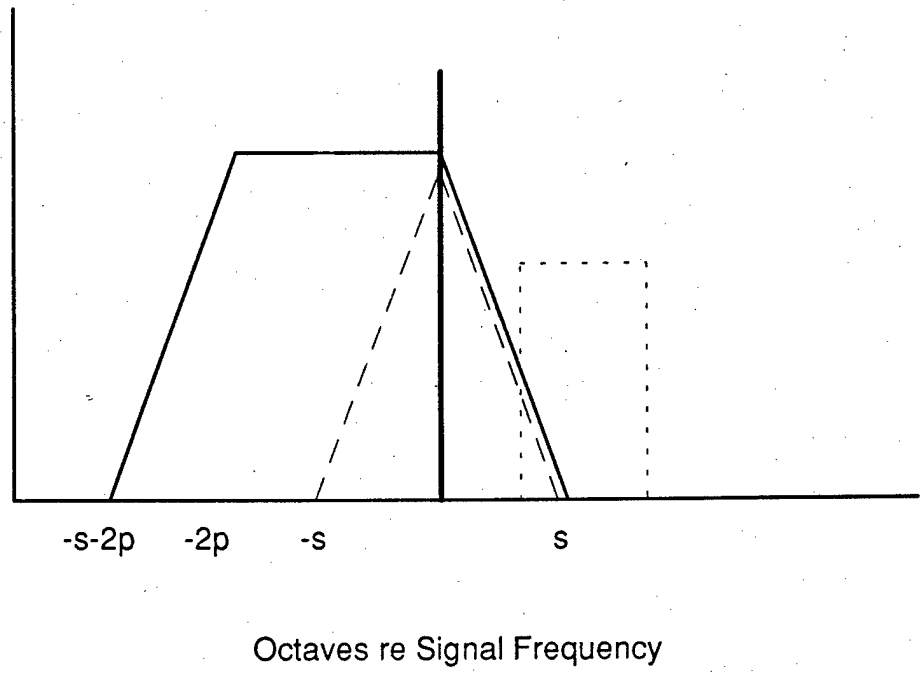

Figure 20. Equivalence between trapezoidal and triangular filter in off-frequency viewing situation. Heavy solid line represents signal, and dotted rectangle represents band of noise. The trapezoidal filter with the highest SNR, represented by light solid line, is positioned so that the signal falls at its elbow. Because all of the stimulus energy falls under the trapezoidal filter's slope, this situation cannot yield an estimate of the extent of the trapezoidal filter's flat portion. A triangular filter, represented by dashed lines (displaced downward for clarity), has no flat portion, and it would yield the same threshold.

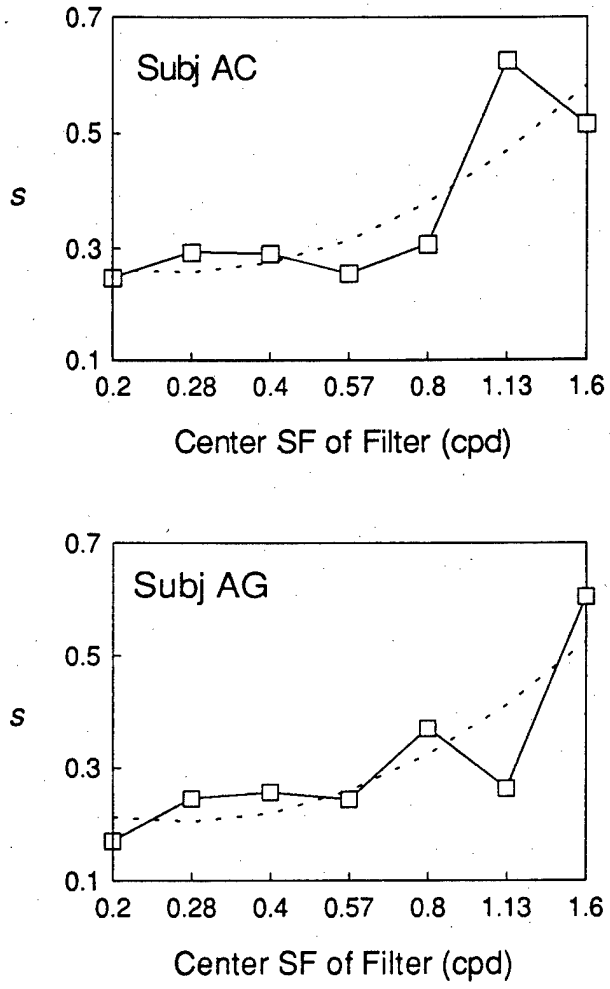

Figure 21. Skirt parameter, *s*, of trapezoidal filter as a function of center SF of filter. Solid line represents values from independent fits at each SF. Dashed line represents best-fitting second-order polynomial in logarithmic SF.

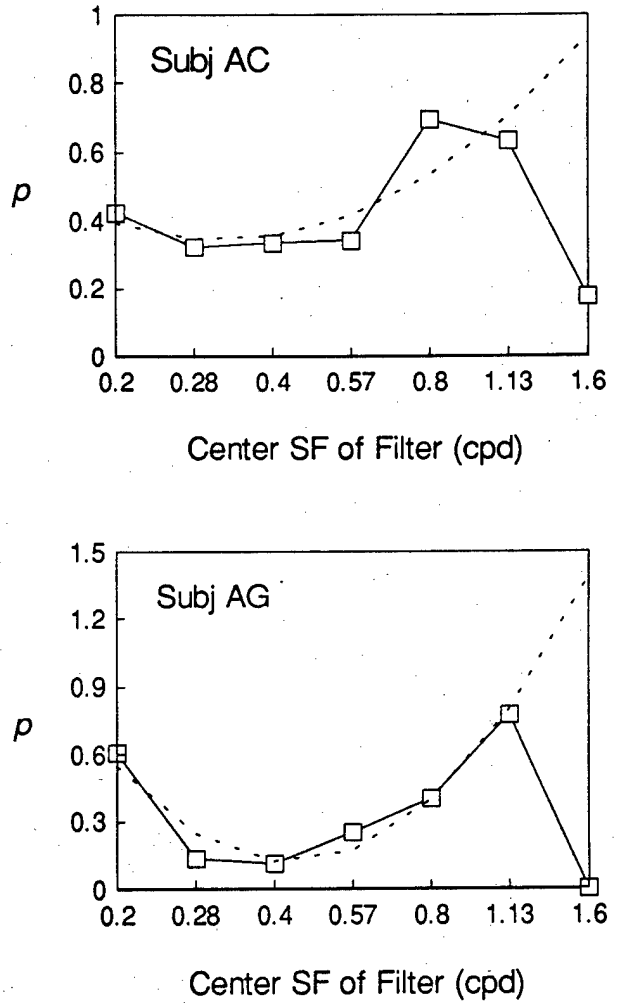

Figure 22. Passband parameter, p, of trapezoidal filter as a function of center SF of filter. Solid line represents values from independent fits at each SF. Dashed line represents best-fitting second-order polynomial in logarithmic SF (fit only based on data through 1.13 cpd).

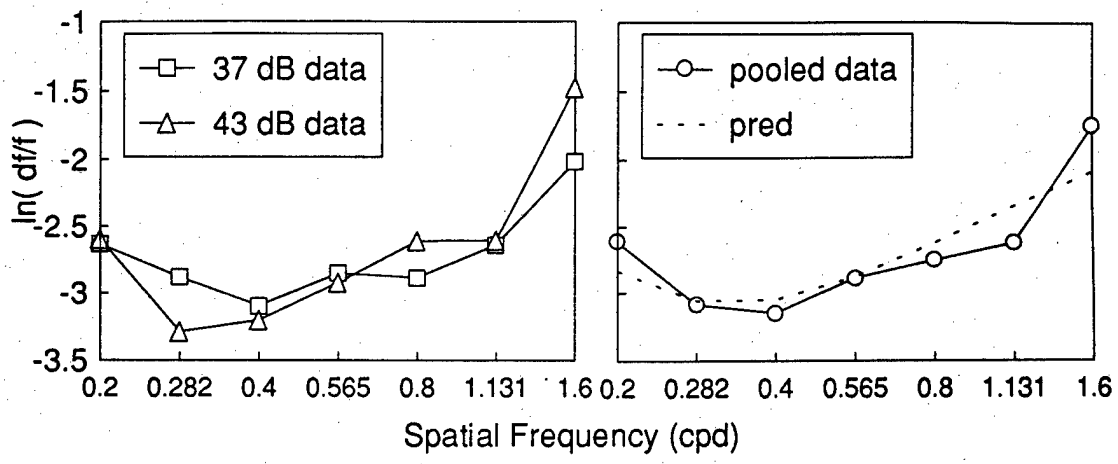

Figure 23. Subject AC's SF discrimination thresholds as a function of SF. Left panel displays data separately for sinusoids with intensities of 37 dB and 43 dB re 0.1 arcmin. Right panel displays data pooled across intensity, with predictions of model.

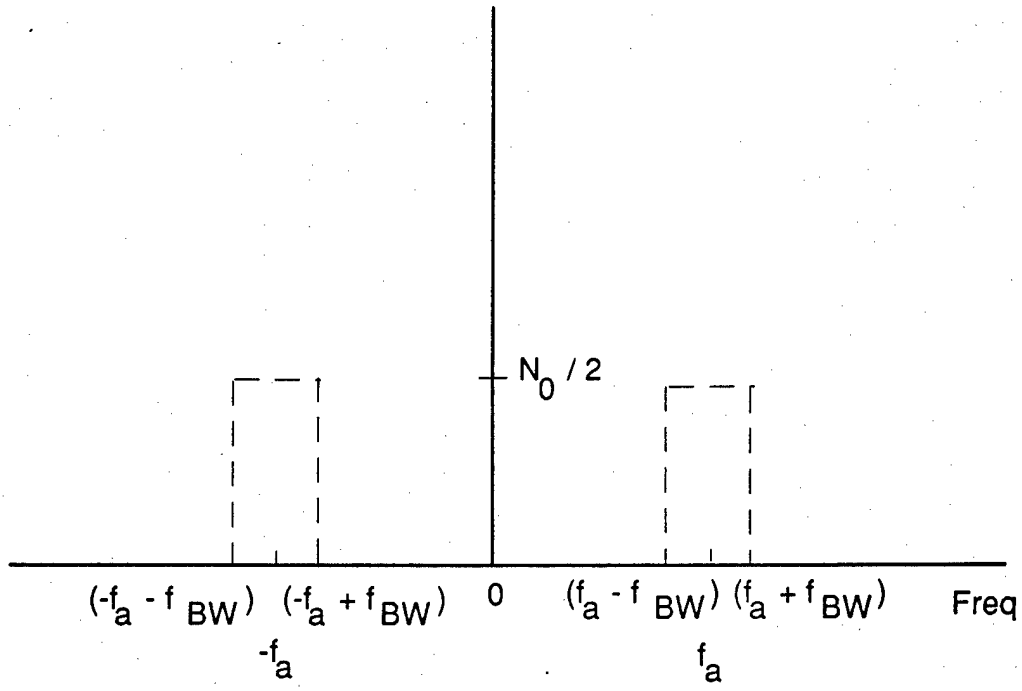

Figure 24. Two-sided power spectrum, $\Gamma(f)$, of narrowband gaussian noise with arithmetic center frequency f_a, arithmetic bandwidth $2f_{BW}$, and spectrum level N_0.

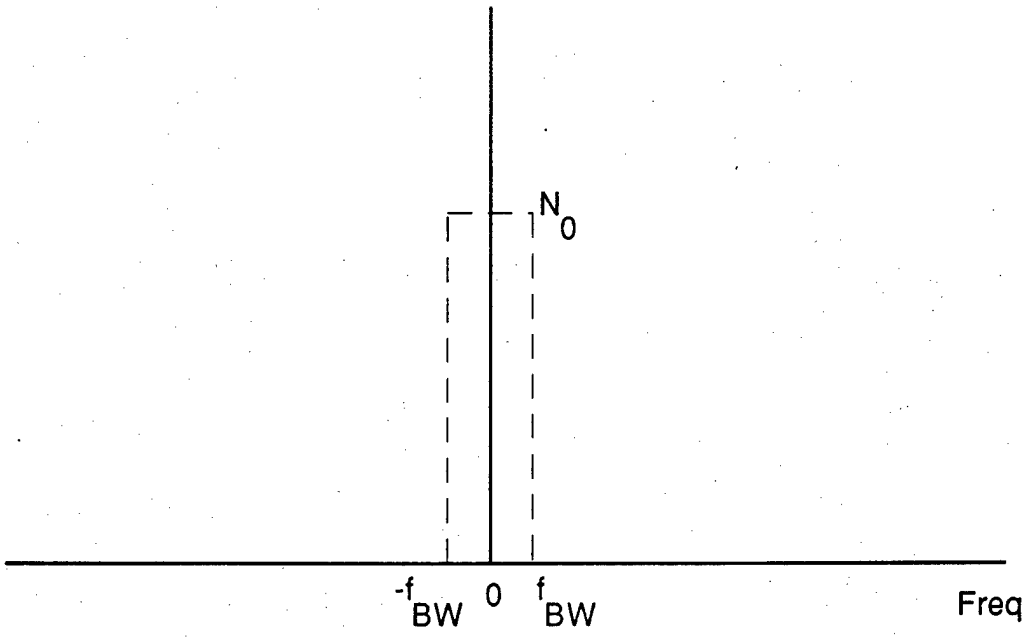

Figure 25. $\tilde{\Gamma}(f)$, a rectangular pulse in frequency.

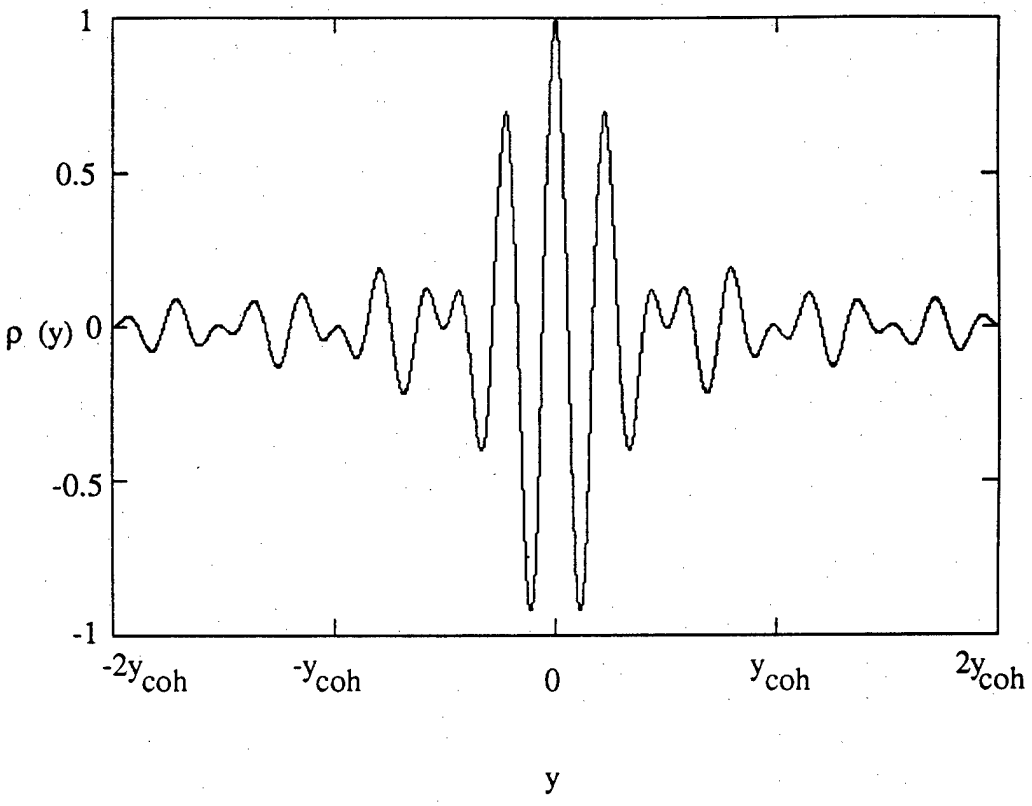

Figure 26. Autocorrelation function of narrowband gaussian noise.

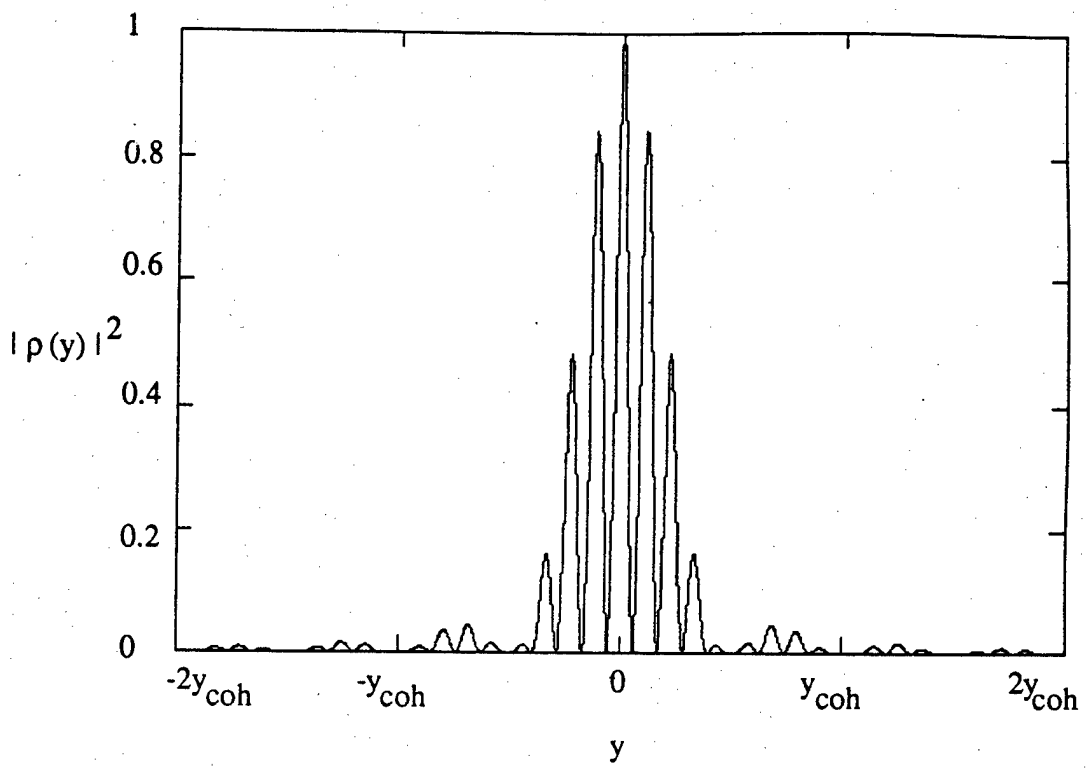

Figure 27. Squared magnitude of autocorrelation function of narrowband gaussian noise.

640 Zamora Ave.
Coral Gables, FL 33134
December 18, 1992

Warren Peele
SCEEE
1101 Massachusetts Ave.
St. Cloud, FL 34769

Dear Prof. Peele:

On December 14, I passed my final oral examination for the degree of Doctor of Philosophy. The one-semester fellowship extension I received was instrumental in the completion of my dissertation, which I have now completed in the last week of the semester.

With the advent of the new year, I will begin a post-doctoral position in the Dept. of Psychology at the University of Miami. My address will be Mailman Center for Child Development, University of Miami, P. O. Box 016820, Miami, FL 33101 (Internet: acobolew@peds.med.miami.edu).

As required according to the terms of my fellowship, I have enclosed a copy of my doctoral dissertation.

Very Truly Yours,

Alan B. Cobo-Lewis, Ph.D.

Enc.